Progress in Mathematics
309

Johannes Kellendonk
Daniel Lenz
Jean Savinien
Editors

Mathematics of Aperiodic Order

Birkhäuser

Progress in Mathematics
Volume 309

Series Editors
Hyman Bass, *University of Michigan, Ann Arbor, USA*
Jiang-Hua Lu, *The University of Hong Kong, Hong Kong SAR, China*
Joseph Oesterlé, *Université Pierre et Marie Curie, Paris, France*
Yuri Tschinkel, *Courant Institute of Mathematical Sciences, New York, USA*

More information about this series at http://www.springer.com/series/4848

Johannes Kellendonk • Daniel Lenz • Jean Savinien
Editors

Mathematics of Aperiodic Order

Birkhäuser

Editors
Johannes Kellendonk
Institut Camille Jordan
Université Claude Bernard Lyon 1
Villeurbanne Cedex, France

Daniel Lenz
Institut für Mathematik
Friedrich-Schiller-Universität Jena
Jena, Germany

Jean Savinien
Institut Elie Cartan de Lorraine
Université de Lorraine
Metz Cedex 1, France

ISSN 0743-1643 ISSN 2296-505X (electronic)
Progress in Mathematics
ISBN 978-3-0348-0902-3 ISBN 978-3-0348-0903-0 (eBook)
DOI 10.1007/978-3-0348-0903-0

Library of Congress Control Number: 2015942704

Mathematics Subject Classification (2010): 52C23, 37B50, 47A35, 11K70, 58B34

Springer Basel Heidelberg New York Dordrecht London
© Springer Basel 2015
This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer Basel AG is part of Springer Science+Business Media (www.birkhauser-science.com)

Contents

Preface vii

1 Non-Periodic Systems with Continuous Diffraction Measures
M. Baake, M. Birkner and U. Grimm 1

2 On the Pisot Substitution Conjecture
S. Akiyama, M. Barge, V. Berthé, J.-Y. Lee and A. Siegel 33

3 Cohomology of Hierarchical Tilings
L. Sadun 73

4 Spaces of Projection Method Patterns and their Cohomology
J. Hunton 105

5 Equicontinuous Factors, Proximality and Ellis Semigroup for Delone Sets
J.-B. Aujogue, M. Barge, J. Kellendonk and D. Lenz 137

6 Linearly Repetitive Delone Sets
J. Aliste-Prieto, D. Coronel, M.I. Cortez, F. Durand and S. Petite ... 195

7 Tilings with Infinite Local Complexity
N. Priebe Frank 223

8 On the Noncommutative Geometry of Tilings
A. Julien, J. Kellendonk and J. Savinien 259

9 Spectral Properties of Schrödinger Operators Arising in the Study of Quasicrystals
D. Damanik, M. Embree and A. Gorodetski 307

10 Additive Properties of Sets and Substitutive Dynamics
S. Puzynina and L.Q. Zamboni 371

11 Delone Sets and Material Science: a Program
J.V. Bellissard 405

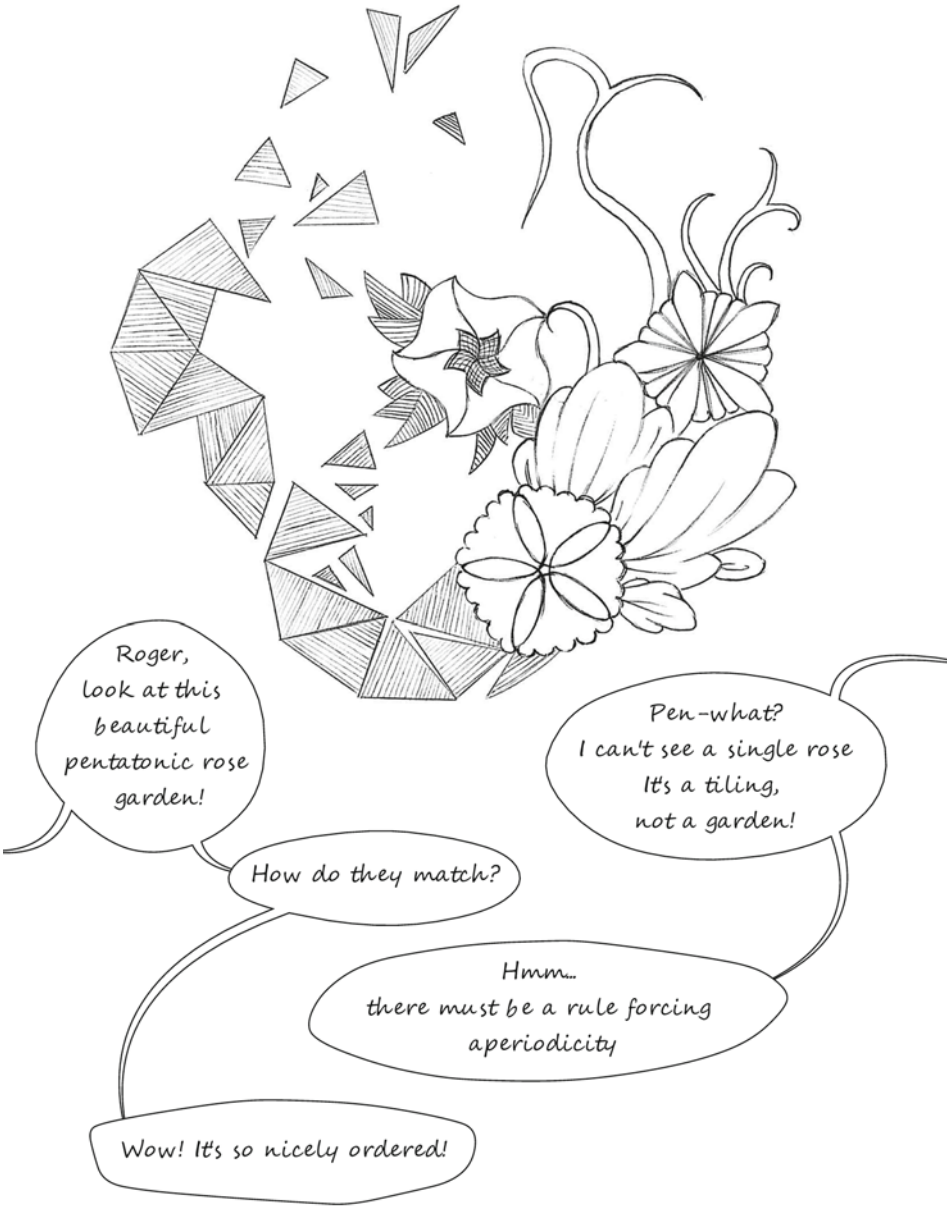
Roger, look at this beautiful pentatonic rose garden!
Pen-what? I can't see a single rose It's a tiling, not a garden!
How do they match?
Hmm... there must be a rule forcing aperiodicity
Wow! It's so nicely ordered!

© 2015 Sarah Kellendonk

Preface

We present here a collection of articles about aperiodic systems which exhibit long range order. By aperiodic systems we mean infinite systems which serve as mathematical models for the description of physical structures which are not the periodic repetition of an "elementary block". Typically aperiodic systems are point patterns in Euclidean space or their functional counterparts, like measures or physical potential functions.

With the discovery of quasicrystals rose the awareness that an aperiodic system may have long range order. Vaguely speaking this means that there is strong coherence between the local structure of the system at different positions in space so that collective phenomena, like diffraction, look very different from those of random systems and rather like those of periodic systems. Ways of making this idea mathematically precise form part of the research program of Aperiodic Order and constitute the topic of this book.

Originally, long range order arose as a concept from X-ray diffraction. The diffraction spectrum of certain materials shows sharp peaks (called Bragg peaks) revealing that light is diffracted by local configurations of atoms which "repeat homogeneously" enough in the material to lead to constructive interferences. This happens for crystals, but not only for them and so a material is considered long range ordered if it shows sharp Bragg peaks under diffraction.

Nowadays one does not base the definition of (long range) order on diffractive properties alone – one does not even insist that the system can be described by a model for which diffraction makes sense – but simply asks: what are the mathematical features which suggest to call an infinite structure ordered, as opposed to random, and how can one quantify the degree of order. In such a context, the notion of order has to be reinvented properly.

In most aperiodic systems the underlying long range order structure is encoded in a point pattern in Euclidean space. Whether these point patterns should be considered as discrete point sets (Delone sets), or tilings, or as discrete measures is not of importance. These are just different approaches to describe the same structure. The point patterns can either be considered as ordered because they are constructed by deterministic rules. The reader will find here in particular hierarchical rules, like substitution rules, or the cut & project method producing quasiperiodic point patterns. Or point patterns are ordered because they satisfy specific mathematical properties, like being pure point diffractive, or highly repetitive, or having finitely generated cohomology groups, or by their dynamical system being close to equicontinuous. The interplay between these two conceptions of order forms part of the theory discussed in this book.

For many years the evolution of the topic has been recorded (and stimulated) by collections of articles rather than monographs. The last such collection, entitled Directions in Mathematical Quasicrystals, dates from 2000. Certainly, our research

has benefited tremendously from these collections. Considerable progress has been obtained in the last fifteen years. This makes us hope that the present collection can serve a similar purpose as earlier works of this type. In this context, it is our great pleasure to point out to the interested reader that very recently the first volume of a monographic introduction with the title Aperiodic Order written by two of the authors contributing to this collection has appeared. This shows clearly that the topic is a most flourishing one and we wish our readers lots of fun with it. We are very grateful to Sarah Kellendonk for here beautiful drawings, which add an artistic touch to the aesthetics of aperiodic tilings.

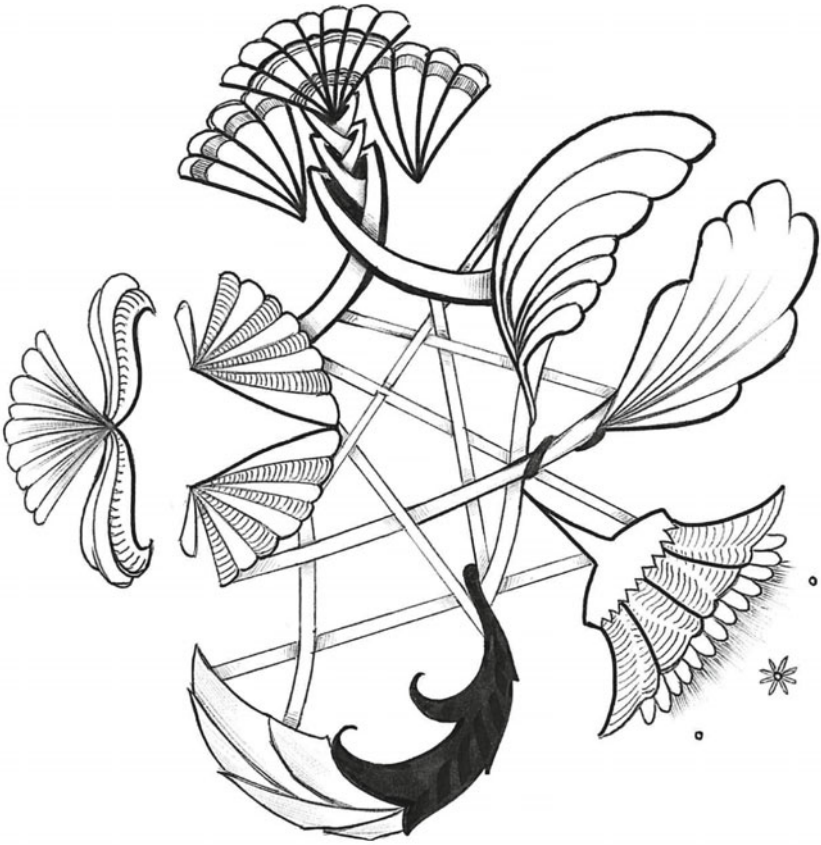

© 2015 Sarah Kellendonk

Short description of the chapters. We start with where it began over 30 years ago: the mathematical theory of diffraction of point sets. It was X-ray diffraction analysis which led to the discovery of quasicrystals and as a consequence to the famous question "Which distributions of matter diffract?", a question with subsequently led to the development of the mathematical theory of diffraction. The **first chapter** presents the new developments in this theory discussing point sets whose diffraction measures have continuous components.

One of the surprising results of the mathematical theory of diffraction is that the diffraction spectrum is related to another kind of spectrum, namely that of the dynamical system associated with the point pattern. In particular, a point pattern is pure point diffractive (its diffraction measure has no continuous part) if and only if the dynamical spectrum is purely discrete (contains only eigenvalues). In that case the group of eigenvalues is generated by the set of Bragg peaks. Therefore, the question as to whether a point pattern is purely diffractive becomes a question about their dynamical system.

For point patterns (or tilings) constructed from substitutions the above question is even of topological nature. There is a famous conjecture in topological dynamics, still open at the time of writing this book, which asks whether point patterns coming from irreducible Pisot substitutions have purely discrete dynamical spectrum. **Chapter two** is devoted to a review about what is known to be true in relation to this "Pisot substitution conjecture". It includes recent attempts to replace the irreducibility assumption by a topological condition.

In the next chapters one of the most fruitful mathematical ideas for the study of point patterns is reviewed, namely that one should not consider a given point pattern on its own but rather the space of all point patterns which look locally like that pattern. This space is called the hull of the pattern, or sometimes the pattern space (or tiling space). It carries a natural topology. This brings in the possibility of applying topological methods to study point patterns. A big runner is cohomology. There are various ways to associate a cohomology group to a point pattern, the quickest to say but least practical being the Čech cohomology of its hull. In **Chapters three** and **four** these various versions of cohomology for point patterns are explained and applied. More specifically, in **Chapter three** the reader will find a discussion of what cohomology is useful for and how it can be computed using the technique of approximating the pattern space by simpler spaces via inverse limits. This is particularly well adapted to compute the cohomology of hierarchical tilings. In **Chapter four** a formulation of pattern cohomology via group cohomology is presented and employed to calculate the cohomology for almost canonical cut & project patterns. In both chapters results from several past years are reviewed but chapter four presents also some background material in detail which so far has only been implicit in the literature, such as the question why the Čech cohomology of a pattern space is the group cohomology of something.

Since our point patterns are subsets of the Euclidean space they may be shifted in the space and this defines the action of the group of translations on the hull of the point pattern. It turns out that the translation action is continuous

and this way we have a topological dynamical system associated with the point pattern. This brings in the possibility of applying methods from topological dynamics systems theory to study point patterns. **Chapter five** is devoted to this. The spectrum of a dynamical system is explained in detail with an emphasis on the role of topological eigenvalues and the maximal equicontinuous factor. These are part of a circle of ideas known in dynamical systems theory for more than 50 years which turn out very useful to study point patterns. In particular a hierarchy of point patterns can be obtained by studying how close their dynamical system is to an equicontinuous one.

The next chapter is devoted to a combinatorial property of point patterns in Euclidean space which is usually considered to describe highly ordered aperiodic materials: linear repetitivity (linear recurrence). In an arbitrarily chosen point pattern a given local configuration might never occur again, or only in one corner of the space. A repetitive point pattern is one where any local configuration appears everywhere again with bounded distance. More precisely, for any radius r there is some (bigger) radius R such that any ball of radius R contains an occurrence (or rather a translated copy) of any local configuration of size r. If there is some constant L such that $R \leq Lr$ the point set is called linearly repetitive. This is the strongest possible case for an aperiodic structure: the point pattern must be periodic if R is bounded by a function of r which grows slower than linearly with r. **Chapter six** discusses the consequences of this property for the underlying dynamical system of point sets. The system is strictly ergodic, never measurably mixing and admits a subadditive ergodic theorem. The continuous and measurable eigenvalues are characterized. It is shown that the dynamics admits only finitely many factors (up to conjugation). Deformation of linearly repetitive point patterns are also studied and related to Lipschitz rectifiability of lattices.

Many of the results discussed in this book hold under the assumption that the point patterns have finite local complexity. This means that, up to translation, there are only finitely many local configurations of a given size. The contribution of **Chapter seven** goes beyond that framework. Methods for the construction of point patterns which have infinite local complexity are discussed, their properties analysed, and many examples given.

Non-commutative geometry has played an early role in the discussion of aperiodic systems, in particular in the attempt to define an analogue of the Brillouin zone for aperiodic solids. The ordinary Brilluoin zone is the dual of the lattice of periods and so shrinks to a point if the lattice of periods is trivial. It has been quickly realised that there exists, however, a non-commutative version in form of a non-commutative C^*-algebra. This has been developped in the 80s and 90s into a fully fledged theory of the non-commutative *topology* of aperiodic systems. Only relatively recently succesful attempts were made to construct the fundamental element of non-commutative *geometry*, namely a spectral triple for aperiodic tilings. So far, convincing constructions of spectral triples have only been obtained for commutative tiling algebras, but they have led to new insight into these tilings. Poles of associated zeta functions are related to the combinatorial complexity of

the tiling, interesting Laplace-type operators may be defined, and a combinatorial property standing for high order (bounded powers) may be characterized by means of these spectral triples. This is reviewed in **Chapter eight**.

A very important aspect of a theory of aperiodic materials is the study of the quantum mechanical motion of a particle in the material. One wishes to solve the Schrödinger equation with a potential created by the sum of atomic potentials for an aperiodic arrangement of atoms, the arrangement being however not random but showing high order. Although this subject has been investigated as early as 1979, the paradigm being that one-dimensional quasiperiodic operators have singular continuous spectrum, it is far from being settled and even relatively crude results need heavy machinery. **Chapter nine** presents a general outline on the type of Schrödinger operators used in quasicrystal physics and then provides an overview on the latest results for one dimensional examples in which the method of trace maps can be successfully applied.

With **Chapter ten** on arithmetic properties of subshifts we include a link between tiling theory and number theory. Arithmetic and combinatorial properties of words are related to dynamical and diffraction properties of associated subshifts. Specifically, this chapter studies the class of central sets and IP-sets of the natural numbers, which have very rich additive properties. In particular it is shown how one can produce a central set as set of occurrences of suitable factors of an infinite word over a finite alphabet with a sufficiently rich combinatorial structure. Moreover, a relation between central sets and the strong coincidence conjecture is discussed (the latter is relevant for the Pisot substitution conjecture, the subject of Chapter two).

The final **Chapter eleven** outlines how a realistic theory of solids or liquid physical systems would have to be based on Delone sets of infinite local complexity. This contribution is a particular one because it is rather a view into the future than a review: it sets out a program to study the dynamics of atoms in solid or liquid phases with the help of Delone sets and their graphs.

February 2, 2015

Johannes Kellendonk
Daniel Lenz
Jean Savinien

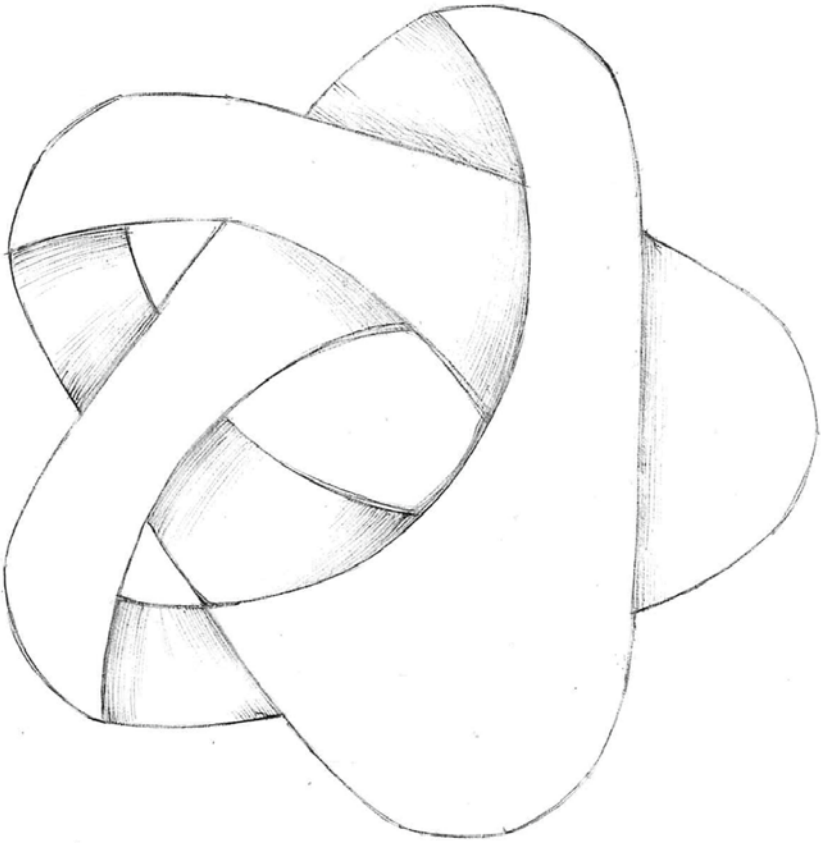

© 2015 Sarah Kellendonk

Progress in Mathematics, Vol. 309, 1–32
© 2015 Springer Basel

Non-Periodic Systems with Continuous Diffraction Measures

Michael Baake, Matthias Birkner and Uwe Grimm

Abstract. The present state of mathematical diffraction theory for systems with continuous spectral components is reviewed and extended. We begin with a discussion of various characteristic examples with singular or absolutely continuous diffraction, and then continue with a more general exposition of a systematic approach via stationary stochastic point processes. Here, the intensity measure of the Palm measure takes the role of the autocorrelation measure in the traditional approach. We furthermore introduce a 'Palm-type' measure for general complex-valued random measures that are stationary and ergodic, and relate its intensity measure to the autocorrelation measure.

Mathematics Subject Classification (2010). Primary 42A38, 37A50; Secondary 37B10, 52C23.

Keywords. Kinematic diffraction, (random) dynamical systems, stationary stochastic processes.

1. Introduction

The (mathematical or kinematic) diffraction theory of systems in Euclidean space with pure point spectrum is rather well understood. Ultimately, this is due to the availability of Poisson's summation formula and its generalisations to the setting of measures (or to tempered distributions); see [12, Sec. 9.2] for a systematic exposition. Beyond results on the spectral nature, this often also provides explicit formulas for the diffraction measure, such as in the cases of lattice-periodic systems and model sets. For these systems, there is also a well-understood connection with the Halmos–von Neumann theorem for the corresponding pure point dynamical spectrum; see [55, 16, 60, 18, 20] for details as well as [12] and references therein for general background.

As soon as one enters the realm of systems with continuous diffraction spectra (or at least with continuous spectral components), the situation changes drastically. As in the case of Schrödinger operator spectra [31], much less is known about

the plethora of possibilities, and there rarely are explicit formulas for the diffraction measures of specific examples. Until recently, explicit results were restricted to simple systems of Bernoulli type (hence with disorder that leads to independent random variables) or to some paradigmatic examples in one dimension (and product systems built from them).

There has now been some progress towards explicitly computable examples in various directions [62, 4, 15]. In particular, both for singular and for absolutely continuous cases, constructive approaches have been more successful than previously anticipated; compare [12, Ch. 10]. Consequently, there is some hope that more systems can be understood in this way. This view is also supported by the recent progress in the understanding of the connection between the dynamical and the diffraction spectrum in this more general situation; see [20] and references therein. At the same time, such examples will improve our intuition about systems with continuous diffraction. Below, this will be reflected by several short sketches of characteristic examples (which are covered in more detail in [12]), before we embark on a more systematic setting via general point process theory. Our focus is on systems in $\mathbb{R}^d$, which is the primary situation to understand, particularly from the applications point of view. Extensions to more general locally compact Abelian groups are possible, but will not be discussed here.

2. Diffraction measures – a brief reminder

Let ω be a locally finite (and possibly complex) measure on $\mathbb{R}^d$, which we primarily view as a linear functional on the space $C_{\mathsf{c}}(\mathbb{R}^d)$ of continuous functions with compact support on $\mathbb{R}^d$, together with some mild extra conditions. In favourable cases, ω will be translation bounded. By the classic Riesz–Markov representation theorem, we may identify the measures defined by this approach with regular Borel measures; for a systematic exposition, we refer to [44, 16] as well as [12, Chs. 8 and 9] and references therein. Particularly important examples comprise the *Dirac measure* δ_x, defined by $\delta_x(g) := g(x)$ for $g \in C_{\mathsf{c}}(\mathbb{R}^d)$, and measures of the form

$$\delta_S := \sum_{x \in S} \delta_x , \tag{2.1}$$

which are known as *Dirac combs*, where $S \subset \mathbb{R}^d$ is uniformly discrete. More generally, we will also consider objects of the form $\sum_{x \in S} w(x)\delta_x$, which can be a measure for a general countable set S, then under suitable conditions on the weight function w. Such measures are referred to as *weighted* Dirac combs.

Recall from [44] or [12] that, if ω is a measure on $\mathbb{R}^d$, the (inverted-conjugate) measure $\widetilde{\omega}$ is defined by $\widetilde{\omega}(g) := \overline{\omega(\widetilde{g})}$ for $g \in C_{\mathsf{c}}(\mathbb{R}^d)$, where $\widetilde{g}(x) := \overline{g(-x)}$. Given a measure ω, consider its *autocorrelation measure*

$$\gamma \,=\, \gamma_\omega := \omega \circledast \widetilde{\omega}, \tag{2.2}$$

where $\circledast$ denotes the volume averaged (or Eberlein) convolution. The latter is defined by

$$\omega \circledast \widetilde{\omega} := \lim_{r\to\infty} \frac{\omega_r * \widetilde{\omega_r}}{\mathrm{vol}(B_r(0))}$$

with $B_r(0)$ the (open) ball of radius r around the origin and $\omega_r := \omega|_{B_r(0)}$. At this stage, we assume the existence of the limit. This will be discussed in more detail later.

If (as in many of our examples) ω is a Dirac comb with lattice support, also γ will be supported on the same lattice (or a subset of it). Concretely, if

$$\omega \,=\, w\,\delta_{\mathbb{Z}} := \sum_{n\in\mathbb{Z}} w(n)\,\delta_n ,$$

with a bounded weight function w say, one finds $\gamma = \eta\delta_{\mathbb{Z}}$ with the positive definite function $\eta\colon \mathbb{Z} \longrightarrow \mathbb{C}$ being defined by

$$\begin{aligned} \eta(m) &:= \lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} w(n)\,\overline{w(n-m)} \\ &= \lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} \overline{w(n)}\, w(n+m), \end{aligned} \tag{2.3}$$

provided that all limits exist. In our exposition below, this existence will follow by suitable applications of Birkhoff's ergodic theorem, applied to the dynamical system of the shift action on the orbit closure of the sequence w or to a similar type of dynamical system; compare [16] for a more general setting. One benefit of this approach will emerge via the Herglotz–Bochner theorem [50].

The autocorrelation measure γ is *positive definite* (or of positive type) by construction, which means that $\gamma(g * \widetilde{g}) \geq 0$ for all $g \in C_{\mathsf{c}}(\mathbb{R}^d)$. It is thus Fourier transformable [27], and the Fourier transform $\widehat{\gamma}$ is a positive measure, called the *diffraction measure* of ω; see [29] for the physics behind this notion, and [44] as well as [12, Ch. 9] for the mathematical theory. Within the framework of kinematic diffraction, it describes the outcome of a scattering experiment by quantifying how much intensity is scattered into a given volume of d-space, and thus is the central object of our interest. By the Lebesgue decomposition theorem, there is a unique splitting

$$\widehat{\gamma} \,=\, \widehat{\gamma}_{\mathsf{pp}} + \widehat{\gamma}_{\mathsf{sc}} + \widehat{\gamma}_{\mathsf{ac}}$$

of the diffraction measure into its *pure point* part $\widehat{\gamma}_{\mathsf{pp}}$, its *singular continuous* part $\widehat{\gamma}_{\mathsf{sc}}$ and its *absolutely continuous* part $\widehat{\gamma}_{\mathsf{ac}}$, with respect to Lebesgue measure λ. The pure point part comprises the 'Bragg peaks' (of which there are at most countably many, so $\widehat{\gamma}_{\mathsf{pp}}$ is a sum over at most countably many Dirac measures with positive weights), while the absolutely continuous part corresponds to the diffuse 'background' scattering which is given by a locally integrable density relative to λ. The singular continuous part is whatever remains – if present, it is a measure

that gives no weight to single points, but is still concentrated to an (uncountable) set of zero Lebesgue measure.

Measures ω which lead to a diffraction $\widehat{\gamma} = \widehat{\gamma}_{\mathsf{pp}}$ are called *pure point diffractive*; examples include lattice-periodic measures and measures based on model sets. These have been studied in detail in the context of diffraction of crystals and quasicrystals; see [9] for a recent review and [12, Chs. 8 and 9] for a systematic exposition. Here, we are concentrating on the other two spectral components, which may also carry important information on the (partial) order which is present in the underlying structure. Pure point spectra are discussed in detail in [23, 16, 17, 18, 12, 59, 57, 58]; for related spectral problems in the context of Schrödinger operators, we refer to [31].

3. Guiding examples

As mentioned above, the understanding of systems with continuous diffraction components is less developed than that of pure point diffractive ones. Still, a better intuition will emerge from a sample of characteristic examples. It is the purpose of this section to provide some of them, while we refer to the literature for further ones [13, 22, 4, 9, 15, 12].

3.1. Thue–Morse sequences

Let us begin with a classic example from the theory of substitution systems that leads to a singular continuous diffraction measure with rather different features in comparison with the Cantor measure, the latter being illustrated in Figure 1. Our example has a long history, which can be extracted from [79, 61, 47, 1]. We confine ourselves to a brief summary of the results, and refer to [7, 12] and references therein for proofs and details.

The classic *Thue–Morse* (TM) sequence can be defined via the one-sided fixed point $v = v_0 v_1 v_2 \ldots$ (with $v_0 = 1$) of the primitive substitution rule

$$\varrho\colon \begin{array}{l} 1 \mapsto 1\bar{1} \\ \bar{1} \mapsto \bar{1}1 \end{array}$$

on the binary alphabet $\{1, \bar{1}\}$. The fixed point is the limit (in the obvious product topology) of the (suitably embedded) iteration sequence

$$\begin{aligned} 1 \stackrel{\varrho}{\longmapsto} 1\bar{1} \stackrel{\varrho}{\longmapsto} 1\bar{1}\bar{1}1 \stackrel{\varrho}{\longmapsto} 1\bar{1}\bar{1}1\bar{1}11\bar{1} \\ \stackrel{\varrho}{\longmapsto} \cdots \longrightarrow v = \varrho(v) = v_0 v_1 v_2 v_3 \ldots \end{aligned}$$

and has a number of distinctive properties [1, 68], for instance

- $v_i = (-1)^{\text{sum of the binary digits of } i}$
- $v_{2i} = v_i$ and $v_{2i+1} = \overline{v_i}$, for all $i \in \mathbb{N}_0$;
- $v = v_0 v_2 v_4 \ldots$ and $\bar{v} = v_1 v_3 v_5 \ldots$
- v is (strongly) cube-free (and hence non-periodic).

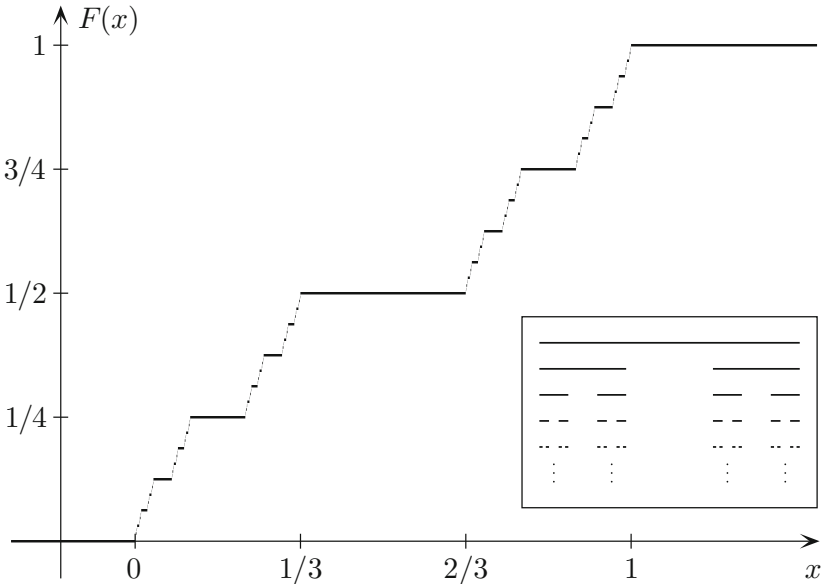

FIGURE 1. The distribution function F of the classic middle-thirds Cantor measure. The construction of the underlying Cantor set is sketched in the inset.

Here, we define $\bar{\bar{1}} = 1$ and identify $\bar{1}$ with -1, also for the later calculations with Dirac combs. A two-sided sequence w can be defined by

$$w(i) \,=\, \begin{cases} v_i, & \text{for } i \geq 0, \\ v_{-i-1}, & \text{for } i < 0, \end{cases}$$

which is a fixed point of ϱ^2, because the seed $w_{-1}|w_0 = 1|1$ is a legal word (it occurs in $\varrho^3(1)$) and $w = \varrho^2(w)$. The (discrete) hull $\mathbb{X} = \mathbb{X}_{\mathrm{TM}}$ of the TM substitution is the closure of the orbit of w under the shift action, which is a subset of $\{\pm 1\}^{\mathbb{Z}}$ and hence a compact space. The orbit of any of its members is dense in $\mathbb{X}$. We thus have a topological dynamical system $(\mathbb{X}, \mathbb{Z})$ that is minimal. When equipped with the standard Borel σ-algebra, the system admits a unique shift-invariant probability measure ν, so that the corresponding measure theoretic dynamical system $(\mathbb{X}, \mathbb{Z}, \nu)$ is strictly ergodic [47, 68].

Any given $w \in \mathbb{X}$ is mapped to a signed Dirac comb ω via

$$\omega \,=\, \sum_{n \in \mathbb{Z}} w(n)\, \delta_n \,.$$

The image of $\mathbb{X}$ is a space of translation bounded measures that is compact in the vague topology. We inherit strict ergodicity via conjugacy, and thus obtain an autocorrelation of the form of Eq. (2.2) with coefficients $\eta(m)$ as in Eq. (2.3). In

fact, this autocorrelation does not depend on the choice of the element from $\mathbb{X}$, so that we may choose the fixed point w from above for the concrete analysis. Due to the nature of w, the coefficients can alternatively be expressed as

$$\eta(m) \,=\, \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} v_n \, v_{n+m}$$

for $m \geq 0$, together with $\eta(-m) = \eta(m)$. It is clear that $\eta(0) = 1$, and the scaling relations of v lead to the recursions [47]

$$\begin{aligned} \eta(2m) \,&=\, \eta(m) \qquad \text{and} \\ \eta(2m{+}1) \,&=\, -\tfrac{1}{2}\bigl(\eta(m) + \eta(m{+}1)\bigr), \end{aligned} \tag{3.1}$$

which are valid for all $m \in \mathbb{Z}$. In particular, the second relation, used with $m = 0$, implies $\eta(1) = -\frac{1}{3}$, which can also be calculated directly.

Since $\eta \colon \mathbb{Z} \longrightarrow \mathbb{C}$ is a positive definite function with $\eta(0) = 1$, there is a unique probability measure μ on the unit circle (which we identify with the unit interval here) such that

$$\eta(m) \,=\, \int_0^1 \mathrm{e}^{2\pi \mathrm{i} m y} \, \mathrm{d}\mu(y)\,, \tag{3.2}$$

which is a consequence of the Herglotz–Bochner theorem [50, Thm. I.7.6]. Since ω is supported on $\mathbb{Z}$, the corresponding diffraction measure $\widehat{\gamma}$ is 1-periodic, which follows from [3, Thm. 1]; see also [12, Sec. 10.3.2]. One then finds the relation

$$\widehat{\gamma} \,=\, \mu * \delta_{\mathbb{Z}}$$

with the measure μ from Eq. (3.2), appropriately interpreted as a measure on $[0, 1)$ and hence also on $\mathbb{R}$. Clearly, one also has $\mu = \widehat{\gamma}|_{[0,1)}$. One can now analyse the spectral type of $\widehat{\gamma}$ via that of the finite measure μ, where we follow [47]; see also [68, 20].

Defining $\Sigma(N) = \sum_{m=-N}^{N} \bigl(\eta(m)\bigr)^2$, a two-step calculation with the recursion (3.1) establishes the inequality $\Sigma(4N) \leq \frac{3}{2}\Sigma(2N)$ for all $N \in \mathbb{N}$. This implies $\lim_{N\to\infty} \Sigma(N)/N = 0$, wherefore Wiener's criterion [79, 50], see also [12, Prop. 8.9], tells us that μ is a continuous measure, so that $\widehat{\gamma}$ cannot have any pure point component. Note that the absence of the 'trivial' pure point component of $\widehat{\gamma}$ on $\mathbb{Z}$ is due to the use of balanced weights, in the sense that 1 and -1 are equally frequent. Consequently, the average weight is zero, and the claim follows from [12, Prop. 9.2].

Let us now define the distribution function F by $F(x) = \mu\bigl([0, x]\bigr)$ for any $x \in [0, 1]$, which is a continuous function that defines a Riemann–Stieltjes measure [53, Ch. X], so that $\mathrm{d}F = \mu$. The recursion relation for η now implies [47] the two functional relations

$$\mathrm{d}F\bigl(\tfrac{x}{2}\bigr) \pm \mathrm{d}F\bigl(\tfrac{x+1}{2}\bigr) \,=\, \bigl\{ {}_{\,-\cos(\pi x)}^{\,1} \bigr\} \, \mathrm{d}F(x)\,,$$

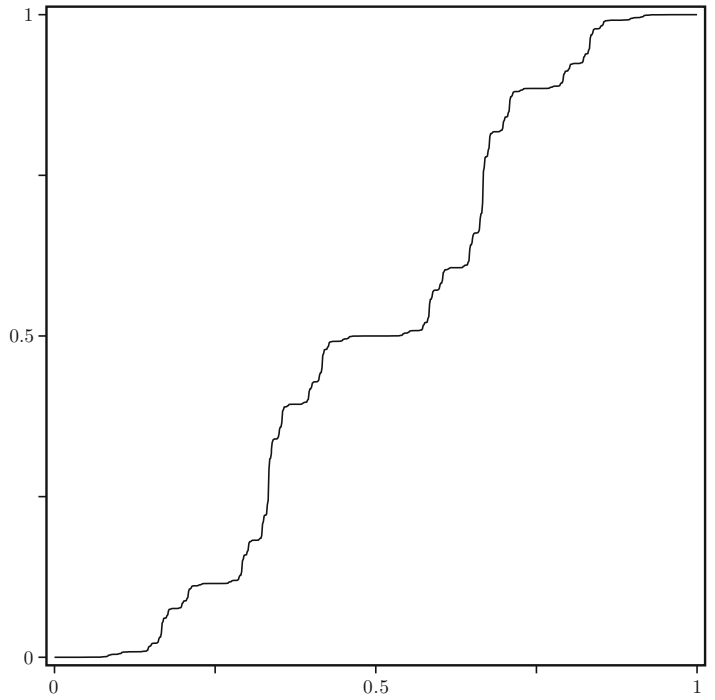

FIGURE 2. The strictly increasing distribution function of the classic, purely singular continuous TM measure on $[0, 1]$.

which have to be satisfied by the ac and sc parts of F separately, because we have $\mu_{\mathsf{ac}} \perp \mu_{\mathsf{sc}}$ in the measure-theoretic sense; see [72, Thm. I.20] or [53, Thm. VII.2.4]. Therefore, defining

$$\eta_{\mathsf{ac}}(m) \, = \int_0^1 \mathrm{e}^{2\pi \mathrm{i} m x} \, \mathrm{d}F_{\mathsf{ac}}(x) \, ,$$

we know that the coefficients $\eta_{\mathsf{ac}}(m)$ must satisfy the same recursions (3.1) as $\eta(m)$, possibly with a different initial condition $\eta_{\mathsf{ac}}(0)$. The classic Riemann–Lebesgue lemma [50, Thm. I.2.8] states that $\lim_{m \to \pm\infty} \eta_{\mathsf{ac}}(m) = 0$. But this limit is only compatible with $\eta_{\mathsf{ac}}(0) = 0$, because $\eta_{\mathsf{ac}}(1) = -\frac{1}{3}\eta_{\mathsf{ac}}(0)$ and $\eta_{\mathsf{ac}}(2m) = \eta_{\mathsf{ac}}(m)$ for all $m \in \mathbb{N}$, so that we must have $\eta_{\mathsf{ac}} \equiv 0$. This means $F_{\mathsf{ac}} = 0$ by the Fourier uniqueness theorem, wherefore μ and hence $\widehat{\gamma}$ (neither of which is the zero measure) are purely singular continuous. The resulting distribution function F is illustrated in Figure 2. Note that F can consistently be extended to a continuous function on $\mathbb{R}$ via $F(x + n) = F(x) + n$ for $n \in \mathbb{Z}$ and then defines $\widehat{\gamma}$ via $\mathrm{d}F = \widehat{\gamma}$ in the Lebesgue–Stieltjes sense. The function F can efficiently be calculated by means of

the uniformly converging Volterra iteration

$$F_{n+1}(x) = \frac{1}{2}\int_0^{2x} \bigl(1-\cos(\pi y)\bigr) F_n'(y)\,\mathrm{d}y \tag{3.3}$$

with $F_0(x) = x$. In contrast to the Devil's staircase of Figure 1, the TM distribution function is *strictly* increasing, which means that there is no plateau (which would indicate a gap in the support of $\widehat{\gamma}$); see [7, 12] and references therein for details and further properties of F. So far, we have obtained the following result.

Theorem 3.1. *Let w be any element of the Thue–Morse hull $\mathbb{X} = \mathbb{X}_{\mathrm{TM}}$, the latter represented as a closed subshift of $\{\pm 1\}^{\mathbb{Z}}$, and consider the corresponding Dirac comb $w\delta_{\mathbb{Z}}$. Then, its autocorrelation γ exists and is given by $\gamma = \eta\delta_{\mathbb{Z}}$ with η being defined by Eq.* (3.1) *together with the initial condition $\eta(0) = 1$.*

The diffraction measure is $\widehat{\gamma} = \mu * \delta_{\mathbb{Z}}$, where μ is the purely singular continuous probability measure from Eq. (3.2)*. In particular, $\widehat{\gamma}$ is purely singular continuous as well.* □

To go one step further, Eq. (3.3) defines an iteration sequence of distribution functions for absolutely continuous measures that converges towards the TM measure in the vague topology. Writing $\mathrm{d}F_n(x) = f_n(x)\,\mathrm{d}x$, one finds

$$f_n(x) \;=\; \prod_{m=0}^{n-1} \bigl(1-\cos(2^{m+1}\pi x)\bigr),$$

which, in the vague limit as $n \to \infty$, gives the well-known Riesz product representation of the TM measure; compare [68] for details and [80] for general background on Riesz products.

The TM sequence is closely related to the limit-periodic *period doubling* (pd) sequence, compare [5, 12] and references therein, via the (continuous) sliding block map defined by

$$\phi\colon \quad 1\bar{1}, \bar{1}1 \mapsto a\,, \quad 11, \bar{1}\bar{1} \mapsto b\,, \tag{3.4}$$

which results in an exact 2-to-1 surjection from the hull $\mathbb{X}_{\mathrm{TM}}$ to $\mathbb{X}_{\mathrm{pd}}$. The latter is the hull of the period doubling substitution defined by

$$\varrho_{\mathrm{pd}}\colon \quad a \mapsto ab\,, \quad b \mapsto aa\,. \tag{3.5}$$

Viewed as topological dynamical systems, this means that $(\mathbb{X}_{\mathrm{pd}}, \mathbb{Z})$ is a factor of $(\mathbb{X}_{\mathrm{TM}}, \mathbb{Z})$. Since both are strictly ergodic, this extends to the corresponding measure-theoretic dynamical systems. The period doubling sequence can be described as a regular model set with a 2-adic internal space [24, 23] and is thus pure point diffractive. This pairing also explains a phenomenon observed in [78], namely that the *dynamical* spectrum of the TM system is richer than its diffraction spectrum. By the dynamical (or von Neumann) spectrum, we mean the spectrum of the unitary operator induced by the shift on the Hilbert space $L^2(\mathbb{X}, \nu)$, where ν is the unique shift-invariant probability measure on $\mathbb{X}$; see [68] for more. Here, the pure point part of the dynamical spectrum is the ring $\mathbb{Z}[\frac{1}{2}]$, which is not even

finitely generated (and only the 'trivial' part $\mathbb{Z}$ is detected by the diffraction measure of the TM system with general weights). In fact, our above measure μ from Theorem 3.1 represents the maximal spectral measure in the ortho-complement of the pure point sector [68, 20]. The missing pure point part, however, is fully recovered via the *diffraction* measure of $\mathbb{X}_{\mathrm{pd}}$; see [12] for details and [20] for a general discussion of this phenomenon.

Various generalisations of this result are known by now. First of all, and perhaps not surprisiningly, this generalises to an entire family of bijective, binary substitutions [5]. Moreover, extensions to higher dimensions are also possible, including the explicit nature of the resulting diffraction measure; compare [40, 11] and references therein.

3.2. Rudin–Shapiro sequence

The (binary) Rudin–Shapiro (RS) chain is a bi-infinite deterministic sequence, with polynomial (in fact linear) complexity function and thus zero entropy. It can be described recursively as $w = (w(n))_{n\in\mathbb{Z}}$ with $w(n) \in \{\pm 1\}$, with initial conditions $w(-1) = -1$, $w(0) = 1$ and the recursion

$$w(4n+\ell) \;=\; \begin{cases} w(n), & \text{for } \ell \in \{0,1\}, \\ (-1)^{n+\ell}\, w(n), & \text{for } \ell \in \{2,3\}, \end{cases} \tag{3.6}$$

which determines $w(n)$ for all $n \in \mathbb{Z}$. The orbit closure of w under the shift action is the (discrete) RS hull $\mathbb{X}_{\mathrm{RS}}$. Alternatively, one can start from a primitive substitution on a 4-letter alphabet (via $a \mapsto ac$, $b \mapsto dc$, $c \mapsto ab$ and $d \mapsto db$) and define a quaternary hull, which then maps to the binary hull via a simple reduction to two letters (for instance via $a, c \mapsto 1$ and $b, d \mapsto -1$); compare [1, 68] or [12, Sec. 4.7.1] for details. The two hulls define topologically conjugate dynamical systems, with local derivation rules in both directions; see [12, Rem. 4.11].

The shift action on $\mathbb{X}_{\mathrm{RS}}$ is strictly ergodic, so that one can define functions $\eta, \vartheta \colon \mathbb{Z} \longrightarrow \mathbb{C}$ via

$$\eta(m) \;=\; \lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} w(n)\, w(n-m) \quad \text{and}$$

$$\vartheta(m) \;=\; \lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} (-1)^n w(n)\, w(n-m),$$

where all limits exist due to unique ergodicity (which is best formulated on the level of the 4-letter alphabet mentioned above). In particular, one finds $\eta(0) = 1$ and $\vartheta(0) = 0$. The recursive structure of Eq. (3.6) now implies the validity of a

closed set of recursive equations [8, 9], namely

$$\begin{aligned}
\eta(4m) &= \tfrac{1+(-1)^m}{2}\,\eta(m),\\
\eta(4m{+}1) &= \tfrac{1-(-1)^m}{4}\,\eta(m) + \tfrac{(-1)^m}{4}\,\vartheta(m) - \tfrac{1}{4}\,\vartheta(m{+}1),\\
\eta(4m{+}2) &= 0,\\
\eta(4m{+}3) &= \tfrac{1+(-1)^m}{4}\,\eta(m{+}1) - \tfrac{(-1)^m}{4}\,\vartheta(m) + \tfrac{1}{4}\vartheta(m{+}1),
\end{aligned}$$

together with

$$\begin{aligned}
\vartheta(4m) &= 0,\\
\vartheta(4m{+}1) &= \tfrac{1-(-1)^m}{4}\,\eta(m) - \tfrac{(-1)^m}{4}\,\vartheta(m) + \tfrac{1}{4}\,\vartheta(m{+}1),\\
\vartheta(4m{+}2) &= \tfrac{(-1)^m}{2}\,\vartheta(m) + \tfrac{1}{2}\,\vartheta(m{+}1),\\
\vartheta(4m{+}3) &= -\tfrac{1+(-1)^m}{4}\,\eta(m{+}1) - \tfrac{(-1)^m}{4}\,\vartheta(m) + \tfrac{1}{4}\,\vartheta(m{+}1),
\end{aligned}$$

which hold for all $m \in \mathbb{Z}$; see [12, Sec. 10.2] for details. A careful inspection shows that the unique solution of this set of equations, with the initial conditions mentioned above, is $\eta(m) = \delta_{m,0}$ together with $\vartheta(m) = 0$ for all $m \in \mathbb{Z}$. Hence, despite the deterministic nature of the RS sequence, the autocorrelation measure is simply given by $\gamma_{\mathrm{RS}} = \delta_0$, so that $\widehat{\gamma_{\mathrm{RS}}} = \lambda$, where λ again denotes Lebesgue measure. Alternatively, the result also follows from the exposition in [68, 67].

Theorem 3.2. *Let w be any element of the Rudin–Shapiro hull $\mathbb{X}_{\mathrm{RS}} \subset \{\pm 1\}^{\mathbb{Z}}$, and consider the corresponding Dirac comb $w\delta_{\mathbb{Z}}$. Then, its autocorrelation exists and is given by $\gamma_{\mathrm{RS}} = \delta_0$, with diffraction measure $\widehat{\gamma_{\mathrm{RS}}} = \lambda$.* □

As in the case of the TM sequence, the non-trivial pure point part of the dynamical spectrum (which is $\mathbb{Z}[\frac{1}{2}]$ once again) is not 'seen' by the diffraction measure, while λ (with multiplicity 2) represents once again the maximal spectral measure in the ortho-complement of the pure point sector. However, the missing pure point component can be recovered by a suitable factor system, the latter obtained via the block map defined by Eq. (3.4). The corresponding factor is represented by a limit-periodic substitution rule that is somewhat reminiscent of the paper folding sequence [1]; see [12, Sec. 10.2] for a complete discussion and [20] for the general connection between dynamical and diffraction spectra. The structure underlying the RS sequence can be generalised to higher-dimensional lattice substitutions in a rather systematic way; see [39] for details.

3.3. Bernoullisation

Let us begin this discussion by recalling the structure of the full Bernoulli shift from the viewpoint of kinematic diffraction. The classic coin tossing process leads to the Dirac comb

$$\omega \,=\, \sum_{n\in\mathbb{Z}} X(n)\,\delta_n\,,$$

where the $(X(n))_{n\in\mathbb{Z}}$ form an i.i.d. family of random variables, each taking values 1 and -1 with probabilities p and $1-p$, respectively. By an application of the strong law of large numbers (SLLN, see [36] for a favourable formulation), almost every realisation has the autocorrelation measure

$$\gamma \,=\, (2p-1)^2\,\delta_{\mathbb{Z}} + 4p(1-p)\,\delta_0\,,$$

and hence (via Fourier transform) the diffraction measure

$$\widehat{\gamma} \,=\, (2p-1)^2\,\delta_{\mathbb{Z}} + 4p(1-p)\,\lambda\,.$$

Here, we have used the classic Poisson summation formula $\widehat{\delta_{\mathbb{Z}}} = \delta_{\mathbb{Z}}$; compare [9] and references therein, as well as [12, Sec. 9.2] for a formulation in the diffraction context. When $p = \frac{1}{2}$, the diffraction boils down to $\widehat{\gamma} = \lambda$. Here, the point part is extinct because the average scattering strength vanishes. For proofs, we refer the reader to [22, 4], while [51, 52] contain several important and non-trivial generalisations and extensions; see also [56] for important related material.

The Bernoulli chain has (metric) entropy [28, 35]

$$h(p) \,=\, -p\log(p) - (1-p)\log(1-p),$$

which is maximal for $p = \frac{1}{2}$, with $h(\frac{1}{2}) = \log(2)$. It vanishes for the deterministic limiting cases $p \in \{0, 1\}$. For the latter, we have $\omega = \mp\delta_{\mathbb{Z}}$, and consequently obtain the pure point diffraction measure $\widehat{\gamma} = \delta_{\mathbb{Z}}$, again via Poisson's summation formula.

Now, the theory of random variables allows for an interpolation between deterministic (binary) sequences and coin tossing sequences as follows. If an element $w \in \{\pm 1\}^{\mathbb{Z}}$ denotes a deterministic sequence (which we assume to be uniquely ergodic for simplicity), consider the random Dirac comb [8]

$$\omega_p \,=\, \sum_{n\in\mathbb{Z}} w(n)X(n)\,\delta_n\,, \tag{3.7}$$

where $(X(n))_{n\in\mathbb{Z}}$ is, as above, an i.i.d. family of random variables with values in $\{\pm 1\}$ and probabilities p and $1-p$. This 'Bernoullisation' of w can be viewed as a 'model of second thoughts', where the sign of the weight at position n is changed with probability $1-p$; compare [12, Sec. 11.2.2].

Let w now be the Rudin–Shapiro sequence from above. By a (slightly more complicated) application of the SLLN, it can be shown [8] that the autocorrelation γ_p of the Dirac comb ω_p is then almost surely given by

$$\gamma_p \,=\, (2p-1)^2\,\gamma_{\mathrm{RS}} + 4p(1-p)\,\delta_0 \,=\, \delta_0\,,$$

irrespective of the value of the parameter $p \in [0, 1]$. Recall that two measures with the same autocorrelation are called *homometric*; see [12, Sec. 9.6] for background. Our observation thus establishes the following classic result; see [8, 9, 12] for details.

Theorem 3.3. *The random Dirac combs ω_p of Eq.* (3.7) *with real parameter values $p \in [0, 1]$ are (almost surely) homometric, with absolutely continuous diffraction measure $\widehat{\gamma_p} = \widehat{\gamma_{\mathrm{RS}}} = \lambda$, irrespective of the value of p. In other words, the family $\{\omega_p \mid p \in [0, 1]\}$ is (almost surely) isospectral.* □

This result shows that diffraction can be insensitive to entropy, because the family of Dirac combs ω_p of Eq. (3.7) continuously interpolates between the deterministic Rudin–Shapiro case with zero entropy and the completely random Bernoulli chain with maximal entropy $\log(2)$. Clearly, the Bernoullisation procedure can be applied to other sequences as well, and can be generalised to higher dimensions. For further aspects of entropy versus diffraction, we refer to [8, 10, 19].

3.4. Random dimers on the line

Another instructive example [25] is based on certain dimer configurations on $\mathbb{Z}$. To formulate it, we follow the exposition in [10] and partition $\mathbb{Z}$ into a close-packed arrangement of 'dimers' (pairs of neighbours), without gaps or overlaps. Clearly, there are just two possibilities to do so, because the position of the first dimer fixes that of all others. Next, decorate each dimer randomly with either $(1,-1)$ or $(-1,1)$, with equal probability. This results in patches such as

$$\cdots[+\ -][-\ +][-\ +][+\ -][-\ +][-\ +][-\ +][+\ -][+\ -]\cdots$$
$$\cdots[-\ +][+\ -][+\ -][-\ +][+\ -][+\ -][+\ -][-\ +][-\ +][+\ -]\cdots$$

where the dimer boxes are indicated by brackets. The set of all decorated sequences defined in this way is given by

$$\mathbb{X} \,=\, \bigl\{w \in \{\pm 1\}^{\mathbb{Z}} \mid M(w) \subset 2\mathbb{Z} \text{ or } M(w) \subset 2\mathbb{Z}+1\bigr\}\,,$$

where $M(w) := \{n \in \mathbb{Z} \mid w(n) = w(n+1)\}$. Note that $M(w)$ is empty precisely for the two periodic sequences that are defined by $w(n) = \pm(-1)^n$ for $n \in \mathbb{Z}$. Clearly, $\mathbb{X} \subset \{\pm 1\}^{\mathbb{Z}}$ is closed and hence compact.

Let $w \in \mathbb{X}$ and consider the corresponding signed Dirac comb on $\mathbb{Z}$ with weights $w(n) \in \{\pm 1\}$. One can then show (again via the SLLN) that the corresponding autocorrelation almost surely exists and is given by [25]

$$\gamma \,=\, \delta_0 - \frac{1}{2}(\delta_1 + \delta_{-1})\,. \tag{3.8}$$

The corresponding diffraction measure is then

$$\widehat{\gamma} \,=\, \bigl(1 - \cos(2\pi k)\bigr)\lambda\,, \tag{3.9}$$

which is again purely absolutely continuous. Here, the (smooth) Radon–Nikodym density relative to λ is written as a function of k. Note that the diffraction measure for general weights h_+ and h_- is given by

$$\widehat{\gamma_{\pm}} \,=\, \frac{|h_+ + h_-|^2}{4}\,\delta_{\mathbb{Z}} + \frac{|h_+ - h_-|^2}{4}\,\widehat{\gamma}$$

with $\widehat{\gamma}$ as in Eq. (3.9). In particular, the measure $\widehat{\gamma_{\pm}}$ shows only the 'trivial' pure point diffraction contribution that arises as the consequence of $\mathbb{Z}$ being the support of the weighted measure under consideration. The same phenomenon also occurs for general (non-balanced) TM and RS sequences; compare [12, Rems. 10.3 and 10.5].

On first sight, the system looks disordered, with entropy $\frac{1}{2}\log(2)$. This seems (qualitatively) reflected by the diffraction. However, the system also defines a

measure-theoretic dynamical system under the action of $\mathbb{Z}$, as generated by the shift. As such, it has a dynamical spectrum that does contain a pure point part, with eigenvalues 0 and $\frac{1}{2}$; we refer to [68] for general background on this concept, and to [25] for the actual calculation of the eigenfunctions. The extension to a (continuous) dynamical system $\mathbb{X}_{\mathrm{c}}$ under the general translation action of $\mathbb{R}$ is done via suspension; see [28, Ch. 11.1] (where the suspension is called a special flow) or [35] for general background.

This finding suggests that some degree of order must be present that is neither visible from the entropy calculation nor from the diffraction measure alone. Indeed, in analogy with the situation of the TM and the RS sequence, one can define a factor of the system by a sliding block map $\phi\colon\, \mathbb{X} \longrightarrow \{\pm 1\}^{\mathbb{Z}}$ which is defined by $(\phi w)(n) = -w(n)w(n+1)$. It maps $\mathbb{X}$ globally 2:1 onto

$$\mathbb{Y} = \phi(\mathbb{X}) = \bigl\{v \in \{\pm 1\}^{\mathbb{Z}} \mid v(n) = 1 \text{ for all } n \in 2\mathbb{Z} \text{ or for all } n \in 2\mathbb{Z}+1\bigr\}.$$

The suspension $\mathbb{Y}_{\mathrm{c}}$ (for the action of $\mathbb{R}$) is defined as above. The mapping ϕ extends accordingly.

The autocorrelation and diffraction measures of the signed Dirac comb $v\delta_{\mathbb{Z}}$ for an element $v \in \mathbb{Y}$ are almost surely given by

$$\gamma \,=\, \frac{1}{2}\delta_0 + \frac{1}{2}\delta_{2\mathbb{Z}} \quad \text{and} \quad \widehat{\gamma} \,=\, \frac{1}{2}\lambda + \frac{1}{4}\delta_{\mathbb{Z}/2}\,.$$

The diffraction of the factor system $\mathbb{Y}$ uncovers the 'hidden' pure point part of the dynamical spectrum, which was absent in the purely absolutely continuous diffraction of the signed Dirac comb $w\delta_{\mathbb{Z}}$ with $w \in \mathbb{X}$. In summary, we have the following situation [25, 20].

Theorem 3.4. *The diffraction measure of the close-packed dimer system $\mathbb{X}$ with balanced weights is purely absolutely continuous and given by Eq.* (3.9)*, which holds almost surely relative to the natural invariant measure of the system.*

The dynamical spectrum of the continuous close-packed dimer system $\mathbb{X}_{\mathrm{c}}$ under the translation action of $\mathbb{R}$ contains the pure point part $\mathbb{Z}/2$ together with a countable Lebesgue spectrum.

The non-trivial part $\mathbb{Z}+\frac{1}{2}$ of the dynamical point spectrum is not reflected by the diffraction spectrum of $\mathbb{X}_{\mathrm{c}}$, but can be recovered via the diffraction spectrum of a suitable factor, such as $\mathbb{Y}_{\mathrm{c}}$. □

As in the case of the Thue–Morse system, where the missing pure point part of the dynamical spectrum is recovered by the diffraction measure of the period doubling factor, we thus see that and how we can recover the missing eigenvalue via a generalised 2-point function. This observation can be extended to symbolic systems over finite alphabets and also to uniquely ergodic Delone dynamical systems of finite local complexity; see [20] for details.

3.5. Ledrappier's shift space

For a long time, people had expected that higher dimensions are perhaps more difficult, but not substantially different. This turned out to be a false premise though, as can be seen from the now classic monograph [73].

In our present context, we pick one characteristic example, the system due to Ledrappier [54], to demonstrate a new phenomenon. We follow the brief exposition in [10] and consider a specific subset of the full shift space $\{\pm 1\}^{\mathbb{Z}^2}$, defined by

$$\mathbb{X}_{\mathrm{L}} \,=\, \bigl\{ w \in \{\pm 1\}^{\mathbb{Z}^2} \mid w(x)\, w(x+e_1)\, w(x+e_2) = 1 \text{ for all } x \in \mathbb{Z}^2 \bigr\}, \qquad (3.10)$$

where e_1 and e_2 denote the standard Euclidean basis vectors in the plane. On top of being a closed subshift, $\mathbb{X}_{\mathrm{L}}$ is also an Abelian group (here written multiplicatively), which then comes with a unique, normalised Haar measure. The latter is also shift invariant, and the most natural measure to be considered in our context; see also the reformulation in terms of Gibbs (or equilibrium) measures in [74].

The system is interesting because the number of patches of a given radius (up to translations) grows exponentially in the radius rather than in the area of the patch. This phenomenon is called *entropy of rank* 1, and indicates a new class of systems in higher dimensions. More precisely, along any lattice direction of $\mathbb{Z}^2$, the linear subsystems essentially behave like one-dimensional Bernoulli chains. It is thus not too surprising that the diffraction measure satisfies the following theorem, though its proof [26] has to take care of the special directions connected with the defining relations of $\mathbb{X}_{\mathrm{L}}$.

Theorem 3.5. *If w is an element of the Ledrappier subshift $\mathbb{X}_{\mathrm{L}}$ of Eq.* (3.10)*, the corresponding weighted Dirac comb $w\delta_{\mathbb{Z}^2}$ has diffraction measure λ, which holds almost surely relative to the Haar measure of $\mathbb{X}_{\mathrm{L}}$.* □

So, the Ledrappier system is homometric to the (full) Bernoulli shift on $\{\pm 1\}^{\mathbb{Z}^2}$, which means that an element of either system almost surely has diffraction measure λ. As mentioned before, via a suitable product of two Rudin–Shapiro chains, also a deterministic system with diffraction λ exists. This clearly demonstrates the insensitivity of pair correlations to the (entropic) type of order or disorder in the underlying system; see also [8]. Due to the defining relation in Eq. (3.10), it is clear that certain three-point correlations in the Ledrappier system cannot vanish, and thus make it distinguishable from the Bernoulli shift.

Although correlation functions of third order can resolve the situation in this case (and in many other examples as well [32, 57]), one can consider other dynamical systems (such as the $(\times 2, \times 3)$-shift [26]) that share almost all correlation functions with the Bernoulli shift on $[0,1]^{\mathbb{Z}^2}$. This is a clear indication that our present understanding of 'order' is incomplete, and that we still lack a good set of tools for the detection and classification of order. For a recent alternative based on direct space statistics, we refer to [6].

3.6. Random matrix ensembles

Another interesting class of random point sets derives from the (scaled) eigenvalue distribution of certain random matrix ensembles; see [14] and references therein. The global eigenvalue distribution of random orthogonal, unitary or symplectic matrix ensembles is known to asymptotically follow the classic semi-circle law. More precisely, this law describes the eigenvalue distribution of the underlying ensembles of symmetric, Hermitian or (symplectically) self-dual matrices with Gaussian distributed entries. The corresponding random matrix ensembles are called GOE, GUE and GSE, with attached β-parameters 1, 2 and 4, respectively. They permit an interpretation as a Coulomb gas, where β is the power in the central potential; see [2, 63] for general background and [34, 38] for the results that are relevant here.

For matrices of dimension N, the semi-circle has radius $\sqrt{2N/\pi}$ and area N. Note that, in comparison with [63], we have rescaled the density by a factor $1/\sqrt{\pi}$, so that we really have a semi-circle (and not a semi-ellipse). To study the local eigenvalue distribution for diffraction, we rescale the central region (between ± 1, say) by $\sqrt{2N/\pi}$. This leads, in the limit as $N \to \infty$, to an ensemble of point sets on the line that can be interpreted as a stationary, ergodic point process of intensity 1; for $\beta = 2$, see [2, Ch. 4.2] and references therein for details. Since the underlying process is simple (meaning that, almost surely, no point is occupied twice), almost all realisations are point sets of density 1.

It is possible to calculate the autocorrelation of these processes, on the basis of Dyson's correlation functions [34]. Though these functions originally apply to the circular ensembles, they have been adapted to the other ensembles by Mehta [63]. For all three ensembles mentioned above, this leads to an autocorrelation of the form

$$\gamma \,=\, \delta_0 + \bigl(1 - f(|x|)\bigr)\lambda \tag{3.11}$$

where f is a locally integrable function that depends on β; see [14] for the explicit formulas, and the left panel of Figure 3 for an illustration.

The diffraction measure is the Fourier transform of γ, which has also been calculated in [34, 63]. Recalling $\widehat{\delta_0} = \lambda$ and $\widehat{\lambda} = \delta_0$, the result is always of the form

$$\widehat{\gamma} \,=\, \delta_0 + \bigl(1 - b(k)\bigr)\lambda \,=\, \delta_0 + h(k)\,\lambda, \tag{3.12}$$

where $b = \widehat{f}$. The Radon–Nikodym density h depends on β and is summarised in [14]. Figure 3 illustrates the result for the three ensembles.

A similar approach is possible on the basis of the eigenvalues of general complex random matrices. This leads to the ensemble studied by Ginibre [63], which is also discussed in [14]. One common feature of the resulting point sets is the effectively repulsive behaviour of the points, which leads to the 'dip' around 0 for $\widehat{\gamma}$. For the two systems mentioned in this section, we omit the formulation of the full results and refer the reader to [14] for details. Further developments around determinantal and related point processes are described in [15].

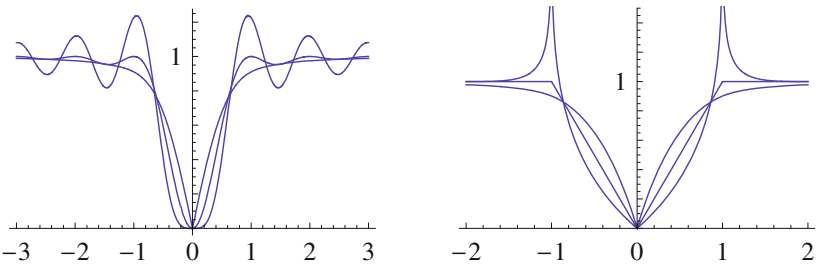

FIGURE 3. Absolutely continuous part of the autocorrelation (left) and the diffraction (right) for the three random matrix derived point set ensembles on the line, with $\beta \in \{1, 2, 4\}$. On the left, the oscillatory behaviour increases with β. On the right, $\beta = 2$ corresponds to the piecewise linear function with bends at 0 and ± 1, while $\beta = 4$ shows a locally integrable singularity at ± 1. The latter reflects the slowly decaying oscillations on the left.

4. The renewal process

A large and interesting class of processes in one dimension can be described as a *renewal* process [37, 4, 14]. Here, one starts from a probability measure μ on $\mathbb{R}_+$ (the positive real line) and considers a machine that moves at constant speed along the real line and drops a point on the line with a waiting time that is distributed according to μ. Whenever this happens, the internal clock is reset and the process resumes. Let us (for simplicity) assume that both the velocity of the machine and the expectation value of μ are 1, so that we end up with realisations that are, almost surely, point sets in $\mathbb{R}$ of density 1 (after we let the starting point of the machine move to $-\infty$, say).

Clearly, the resulting process is stationary and can thus be analysed by considering all realisations which contain the origin. Moreover, there is a clear (distributional) symmetry around the origin, so that we can determine the corresponding autocorrelation γ of almost all realisations from studying what happens to the right of 0. Indeed, if we want to know the frequency per unit length of the occurrence of two points at distance x (or the corresponding density), we need to sum the contributions that x is the first point after 0, the second point, the third, and so on. In other words, we almost surely obtain the autocorrelation

$$\gamma \,=\, \delta_0 + \nu + \widetilde{\nu} \tag{4.1}$$

with $\nu = \mu + \mu * \mu + \mu * \mu * \mu + \cdots$, where the proper convergence of the sum of iterated convolutions follows from [4, Lemma 4] or from [12, Sec. 11.3]. Note that the point measure at 0 simply reflects the fact that the almost sure density of the resulting point set is 1. Indeed, ν is a translation bounded positive measure, and satisfies the renewal relations (compare [37, Ch. XI.9] or [4, Prop. 1] for a proof)

$$\nu \,=\, \mu + \mu * \nu \qquad \text{and} \qquad (1 - \widehat{\mu}\,)\,\widehat{\nu} \,=\, \widehat{\mu}\,, \tag{4.2}$$

where $\widehat{\mu}$ is a uniformly continuous and bounded function on $\mathbb{R}$. The second equation emerges from the first by Fourier transform, but has been rearranged to highlight the relevance of the set $S = \{k \mid \widehat{\mu}(k) = 1\}$ of singularities. In this setting, the measure γ of Eq. (4.1) is both positive and positive definite.

Based on the structure of the support of the underlying probability measure μ, one can determine the diffraction of the renewal process explicitly. To do so for a probability measure μ on $\mathbb{R}_+$ with mean 1, we assume the existence of a moment of μ of order $1+\varepsilon$ for some $\varepsilon > 0$; we refer to [4] for details on this condition. The diffraction measure of the point set realisations of the stationary renewal process based on μ almost surely is of the form

$$\widehat{\gamma} \,=\, \widehat{\gamma}_{\mathsf{pp}} + (1-h)\,\lambda,$$

where h is a locally integrable function on $\mathbb{R}$ that is continuous almost everywhere. The pure point part is trivial, meaning $\widehat{\gamma} = \delta_0$, unless the support of μ is contained in a lattice. The details are stated below in Theorem 4.1. Proofs of these claims as well as further results can be found in [4, 14, 12].

The renewal process is a versatile method to produce interesting point sets on the line. These include random tilings with finitely many intervals (which are Delone sets) as well as the homogeneous Poisson process on the line (where μ is the exponential distribution with mean 1); see [4, Sec. 3] for explicit examples and applications. In particular, if one employs a suitably normalised version of the Gamma distribution, one can formulate a one-parameter family of renewal processes that continuously interpolates between the Poisson process (total positional randomness) and the lattice $\mathbb{Z}$ (perfect periodic order); see [4, Ex. 3] for more. The general result reads as follows.

Theorem 4.1. *Let ϱ be a probability measure on $\mathbb{R}_+$ with mean 1, and assume that a moment of ϱ of order $1+\varepsilon$ exists for some $\varepsilon > 0$. Then, the point sets obtained from the stationary renewal process based on ϱ almost surely have a diffraction measure of the form*

$$\widehat{\gamma} \,=\, \widehat{\gamma}_{\mathsf{pp}} + (1-h)\,\lambda\,,$$

where h is a locally integrable function on $\mathbb{R}$ that is continuous except for at most countably many points (namely those of the set $S = \{k \mid \widehat{\varrho}(k) = 1\}$). On $\mathbb{R} \setminus S$, the function h is given by

$$h(k) \,=\, \frac{2\left(|\widehat{\varrho}(k)|^2 - \mathrm{Re}(\widehat{\varrho}(k))\right)}{|1-\widehat{\varrho}(k)|^2}\,.$$

Moreover, the pure point part is

$$\widehat{\gamma}_{\mathsf{pp}} \,=\, \begin{cases} \delta_0, & \text{if } \mathrm{supp}(\varrho) \text{ is not a subset of a lattice,} \\ \delta_{\mathbb{Z}/b}, & \text{otherwise,} \end{cases}$$

where $b\mathbb{Z}$ is the coarsest lattice that contains $\mathrm{supp}(\varrho)$. □

In one dimension, the renewal process allows an efficient derivation of the diffraction of random tilings, which we briefly summarise now.

5. Random tilings

The deterministic Fibonacci chain can be defined by the primitive substitution rule $a \mapsto ab$, $b \mapsto a$, which defines a strictly ergodic (discrete) hull. When a and b are replaced by intervals of length $\tau = \frac{1}{2}(1+\sqrt{5})$ and 1, respectively, the left endpoints of the intervals define a model set (or cut and project set). The corresponding Dirac comb leads to the pure point diffraction measure

$$\widehat{\gamma_{\mathrm{F}}} = \sum_{k \in \frac{1}{\sqrt{5}}\mathbb{Z}[\tau]} I(k)\, \delta_k$$

with intensities $I(k) = \bigl(\frac{\tau}{\sqrt{5}}\, \frac{\sin(\pi\tau k')}{\pi\tau k'}\bigr)^2$. Here, $\frac{\tau}{\sqrt{5}} = \frac{\tau+2}{5}$ is the density of the point set, and k' denotes the algebraic conjugate of k, which is defined on the field $\mathbb{Q}(\sqrt{5})$ by $\sqrt{5} \mapsto -\sqrt{5}$ and acts as the $\star$-map for the underlying model set description. In particular, the diffraction is the same for all Dirac combs of the Fibonacci hull; see [12, Sec. 9.4.1] and references therein for details. An illustration is shown in the upper panel of Figure 4.

The corresponding random tiling ensemble consists of all tilings of the real line by the two types of intervals. For a direct comparison, it makes more sense to only consider those tilings with the same relative frequency of interval types, which means frequencies $1/\tau$ and $1/\tau^2$ for the long and the short interval, respectively.

The diffraction of a typical Dirac comb out of this class was originally derived in [13], but can also be obtained via an application of the renewal structure from Theorem 4.1. This leads to

$$\widehat{\gamma_{\mathrm{rt}}} = \frac{\tau+1}{5}\delta_0 + h\lambda$$

with the Radon–Nikodym density

$$h(k) = \frac{\tau+2}{5}\, \frac{(\sin(\pi k/\tau))^2}{\tau^2(\sin(\pi k\tau))^2 + \tau(\sin(\pi k))^2 - (\sin(\pi k/\tau))^2}.$$

Except for the trivial Bragg peak at $k=0$, the diffraction measure is thus absolutely continuous. Still, the resemblance between this function and the diffraction of the perfect Fibonacci chain is remarkable, as can be seen from Figure 4.

The situation in dimensions $d \geq 2$ is less favourable from a mathematical perspective, although one has a rather clear intuition of what one should expect [43, 70], based on solid scaling arguments. In dimensions $d \geq 3$, a mixed spectrum with pure point and absolutely continuous components is conjectured, while $d=2$ is the critical dimension in the sense that random tilings with non-crystallographic symmetries should display a singular continuous component; see [12, Sec. 11.6.2] for an example.

Unfortunately, only few results have been proved so far. Among them are a rigorous treatment of planar random tiling ensembles with crystallographic symmetries (such as the lozenge tiling and several relatives, see [13, 45, 46]), a group-theoretic approach to one of the random tiling hypotheses [69, 70] and a treatment

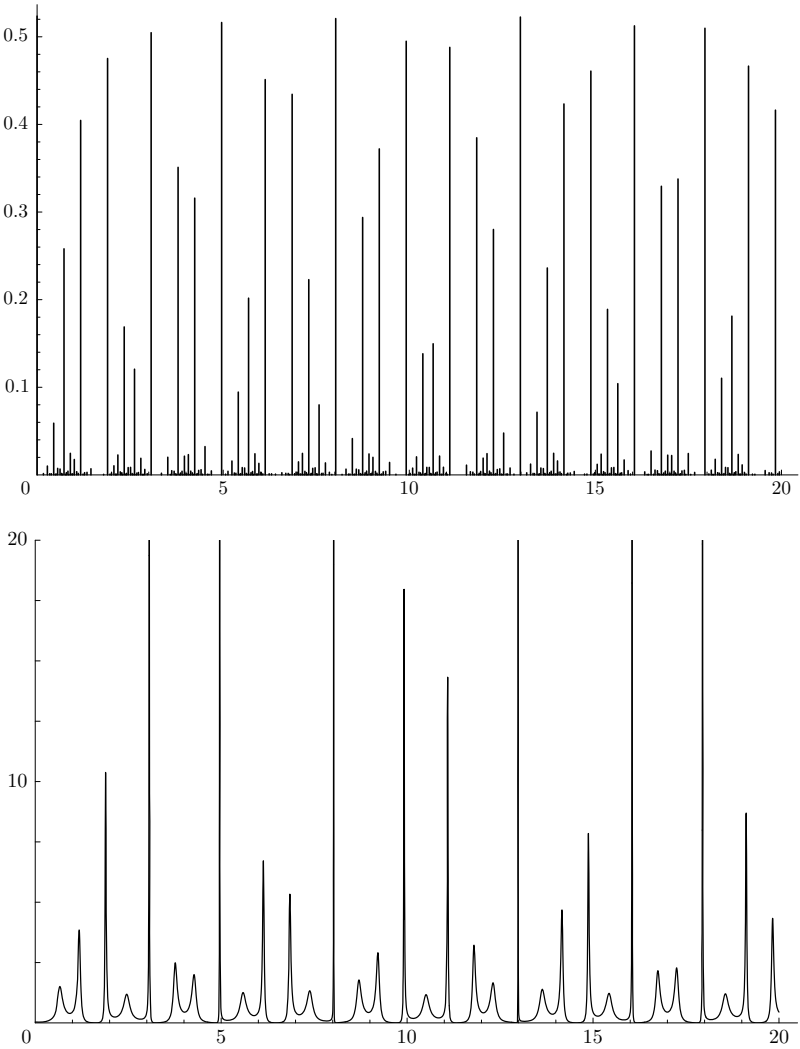

FIGURE 4. The pure point diffraction measure of the perfect Fibonacci chain (upper panel) and the absolutely continuous part of the corresponding random tiling (lower panel). Bragg peaks (in the upper picture) are shown as lines, where the height is the intensity, while the smooth Radon–Nikodym density in the lower picture is truncated at a value of 20 to illustrate the spikyness. The central peak (of intensity $\frac{\tau+1}{5}$) is omitted in both diagrams.

of dense Dirac combs with pure point diffraction [71, 59] that is needed to understand the pure point part of the random tiling diffraction in dimensions $d \geq 3$. The remaining questions are still open, though there is little doubt that the original analysis from [43] is essentially correct.

Let us now leave the realm of explicit examples and turn our attention to a more general approach of systems with randomness, formulated with methods from the theory of point processes; compare [42, 57, 58, 4] for related aproaches and results.

6. Stochastic point processes and the Palm measure

In this section, we take the viewpoint of a general shift-invariant random measure and relate its realisation-wise diffraction to its second moment measure. As such, this section is a complex-valued extension of [4, Sec. 5].

Let $\mu = \mu_{\Re} + \mathrm{i}\mu_{\Im}$ be a locally finite complex-valued measure on $\mathbb{R}^d$ (which means that $\mu_{\Re}$ and $\mu_{\Im}$ are both locally finite signed measures). A short calculation reveals that, for $f \in C_{\mathsf{c}}(\mathbb{R}^d, \mathbb{C})$ of the form $f = g + \mathrm{i}h$ with real-valued g and h, the measure $\widetilde{\mu}$ can consistently be defined via

$$\widetilde{\mu}(f) := \overline{\mu(\widetilde{f})} = \overline{\mu}(f_-) = \mu_{\Re}(f_-) - \mathrm{i}\mu_{\Im}(f_-),$$

where $\widetilde{f}(x) = \overline{f_-(x)}$ with $f_-(x) = f(-x)$. In particular, note that

$$\overline{\mu} = \mu_{\Re} - \mathrm{i}\mu_{\Im} \quad \text{and} \quad \overline{\mu(f)} = \overline{\mu}(\overline{f})$$

hold as expected. The point here is that, after having dealt with the case of real (or signed) measures, the extension to complex measures is canonical and consistent.

To continue, recall the *polar representation* of a complex measure from [33, Ch. XIII.16]; see also [12, Prop. 8.3]. Given μ, there is a measurable function $\alpha_\mu \colon \mathbb{R}^d \to [0, 2\pi)$ such that, for $f \in C_{\mathsf{c}}(\mathbb{R}^d, \mathbb{C})$, one has

$$\int_{\mathbb{R}^d} f(x)\, \mathrm{d}\mu(x) = \int_{\mathbb{R}^d} f(x)\, \mathrm{e}^{\mathrm{i}\alpha_\mu(x)}\, \mathrm{d}|\mu|(x),$$

where $|\mu|$ is the total variation measure of μ. This means that $|\mu|$ is the smallest non-negative measure such that $|\mu(A)| \leq |\mu|(A)$ for any bounded and measurable A, where $|\mu| \leq |\mu_{\Re}| + |\mu_{\Im}|$; compare [12, Sec. 8.5.1] and references therein.

Let $\mathcal{M}$ denote the $\mathbb{C}$-vector space of all locally finite, complex-valued measures ϕ on $\mathbb{R}^d$, so $\phi \in \mathcal{M}$ means $|\phi(A)| < \infty$ for any bounded Borel set A. A sequence $(\phi_n)_{n\in\mathbb{N}} \subset \mathcal{M}$ converges vaguely to ϕ if $\phi_n(f) \longrightarrow \phi(f)$ as $n \to \infty$ for all $f \in C_{\mathsf{c}}(\mathbb{R}^d)$. The space $\mathcal{M}$ is closed in the topology of vague convergence of measures (in fact, $\mathcal{M}$ is a Polish space with this topology). We let $\Sigma_{\mathcal{M}}$ denote the σ-algebra of Borel sets of $\mathcal{M}$. The latter can be described as the σ-algebra of subsets of $\mathcal{M}$ generated by the requirement that, for all bounded Borel sets $A \subset \mathbb{R}^d$, the mapping $\phi \mapsto \phi(A)$ is measurable.

For each $t \in \mathbb{R}^d$, let T_t denote the translation operator on $\mathbb{R}^d$, as defined by the mapping $x \mapsto t + x$. Clearly, $T_t T_s = T_{t+s}$, and the inverse of T_t is given by $T_t^{-1} = T_{-t}$. For functions f on $\mathbb{R}^d$, the corresponding translation action is defined via $T_t f = f \circ T_{-t}$, so that $(T_t f)(x) = f(x-t)$. Similarly, for $\phi \in \mathcal{M}$, let $T_x\phi := \phi \circ T_{-x}$ be the image measure under the translation, so that $(T_x\phi)(A) = \phi(T_{-x}(A)) = \phi(A-x)$ for any measurable subset $A \subset \mathbb{R}^d$, and $(T_x\phi)(f) = \int_{\mathbb{R}^d} f(y)\,\mathrm{d}(T_x\phi)(y) = \int_{\mathbb{R}^d} f(x+z)\,\mathrm{d}\phi(z) = \phi(T_{-x}f)$ for functions. This means that there is a translation action of $\mathbb{R}^d$ on $\mathcal{M}$. Finally, we also have a translation action on $\mathcal{P}(\mathcal{M})$, the probability measures on $\mathcal{M}$, via $(T_x Q)(A) = Q(T_{-x}A)$ for any $A \in \Sigma_{\mathcal{M}}$ and $Q \in \mathcal{P}(\mathcal{M})$. A set $A \in \Sigma_{\mathcal{M}}$ is called *invariant* (under translations) if $T_{-x}A = A$ for all $x \in \mathbb{R}^d$.

A (complex-valued) *random measure* $\boldsymbol{\Phi}$ is a random variable (defined on some probability space $(\Theta, \mathcal{F}, \pi)$) with values in $\mathcal{M}$, which formally means that $\boldsymbol{\Phi} : \Theta \longrightarrow \mathcal{M}$ is an $(\mathcal{F} - \Sigma_{\mathcal{M}})$-measurable function. Its distribution is then $Q = \pi \circ \boldsymbol{\Phi}^{-1} \in \mathcal{P}(\mathcal{M})$, i.e., the image measure of π under $\boldsymbol{\Phi}$. We will follow the usual practice in probability theory and not make the underlying probability space explicit (a canonical choice can in many cases simply be $\Theta = \mathcal{M}$ and $\boldsymbol{\Phi} = \mathrm{Id}_{\mathcal{M}}$). We will also usually suppress the dependence of $\boldsymbol{\Phi}$ on $\theta \in \Theta$ in the notation. Integrals over Θ w.r.t. the probability measure π will be denoted by $\mathbb{E}$, the expectation value.

$\boldsymbol{\Phi}$ is called *stationary* if its distribution Q satisfies $T_x Q = Q$ for all $x \in \mathbb{R}^d$. A stationary random measure is called *ergodic* if the shift-invariant σ-algebra is trivial, which means that any invariant A has probability 0 or 1 (more generally, one requires $Q(A) \in \{0,1\}$ whenever $Q\big((T_{-x}A) \triangle A\big) = 0$ for all $x \in \mathbb{R}^d$; compare [30, Def. 10.3.I and Prop. 10.3.III]).

In what follows, we generally assume that

$$\begin{minipage}{0.8\textwidth}$\boldsymbol{\Phi}$ is a (possibly) complex-valued, stationary and ergodic random measure on $\mathbb{R}^d$,\end{minipage} \tag{6.1}$$

which means that there is a decomposition $\boldsymbol{\Phi} = \boldsymbol{\Phi}_{\Re} + \mathrm{i}\boldsymbol{\Phi}_{\Im}$ where both $\boldsymbol{\Phi}_{\Re}$ and $\boldsymbol{\Phi}_{\Im}$ are signed, real-valued, stationary, ergodic random measures on $\mathbb{R}^d$. To verify the last statement, note that, since for any bounded measurable $A \subset \mathbb{R}^d$, the mapping $\theta \mapsto \boldsymbol{\Phi}(A)\,(= \boldsymbol{\Phi}(\theta, A)) \in \mathbb{C}$ is measurable, also $\boldsymbol{\Phi}_{\Re}(A)$ and $\boldsymbol{\Phi}_{\Im}(A)$ are measurable as functions of θ. Consider any shift-invariant measurable $B \subset \mathcal{M}_{\mathrm{real}}$ ($\mathcal{M}_{\mathrm{real}}$ denotes the locally finite signed measures on $\mathbb{R}^d$), then $\{\boldsymbol{\Phi} \mid \boldsymbol{\Phi}_{\Re} \in B\}$ is shift invariant and measurable as well, so $\mathbb{P}(\boldsymbol{\Phi}_{\Re} \in B) \in \{0,1\}$, and analogously for $\boldsymbol{\Phi}_{\Im}$. We further assume that $\boldsymbol{\Phi}$ is locally square integrable in the sense that

$$\mathbb{E}\left[\big(|\boldsymbol{\Phi}_{\Re}|(A)\big)^2 + \big(|\boldsymbol{\Phi}_{\Im}|(A)\big)^2\right] < \infty \qquad \text{for all bounded } A \subset \mathbb{R}^d, \tag{6.2}$$

where $|\boldsymbol{\Phi}_{\Re}|$ and $|\boldsymbol{\Phi}_{\Im}|$ denote the total variation measures of $\boldsymbol{\Phi}_{\Re}$ and $\boldsymbol{\Phi}_{\Im}$, respectively.

In analogy with the real-valued case in [4, Sec. 5.2], we define $\mu^{(2)}$, the second moment measure of $\varPhi$, via

$$\mu^{(2)}(A \times A') \,=\, \mathbb{E}\big[\varPhi(A)\overline{\varPhi(A')}\big] \qquad \text{for bounded } \; A, A' \in \mathcal{B}(\mathbb{R}^d), \tag{6.3}$$

hence, for $f \in C_{\mathsf{c}}(\mathbb{R}^d \times \mathbb{R}^d, \mathbb{C})$,

$$\mu^{(2)}(f) \,=\, \mathbb{E}\left[\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(x,y)\,\mathrm{d}\varPhi(x)\,\mathrm{d}\overline{\varPhi}(y)\right].$$

By the shift invariance of the distribution of $\varPhi$, we have

$$\int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2)}(x,y) \,=\, \int_{\mathbb{R}^d\times\mathbb{R}^d} f(x+z,y+z)\,\mathrm{d}\mu^{(2)}(x,y)$$

for all $z \in \mathbb{R}^d$, and hence we can factor out this symmetry to obtain the reduced second moment measure $\mu^{(2)}_{\mathrm{red}}$. The latter is a locally finite complex-valued measure that is characterised by

$$\int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2)}(x,y) \,=\, \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u+v,u)\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(v)\,\mathrm{d}\lambda(u) \tag{6.4}$$

for $f \in C_{\mathsf{c}}(\mathbb{R}^d \times \mathbb{R}^d, \mathbb{C})$. By the shift invariance of Lebesgue measure on $\mathbb{R}^d$, we equivalently have

$$\int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2)}(x,y) \,=\, \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u,u-v)\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(v)\,\mathrm{d}\lambda(u). \tag{6.5}$$

To prove the existence of $\mu^{(2)}_{\mathrm{red}}$, one can decompose $\mu^{(2)} = \mu^{(2)}_{\Re} + \mathrm{i}\mu^{(2)}_{\Im}$ into real and imaginary parts and then use the well-known real-valued results (compare [30, Lemma 10.4.III]) to obtain $\mu^{(2)}_{\mathrm{red}} = \mu^{(2)}_{\Re,\mathrm{red}} + \mathrm{i}\mu^{(2)}_{\Im,\mathrm{red}}$.

Note that $\mu^{(2)}_{\mathrm{red}}$ is uniquely defined and is a positive definite measure, since

$$\begin{aligned}
\mu^{(2)}\big(f \otimes \overline{g}\big) \,&=\, \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u+v)\,\overline{g(u)}\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(v)\,\mathrm{d}\lambda(u)\\
&=\, \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(v-w)\,\overline{g(-w)}\,\mathrm{d}\lambda(w)\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(v)\\
&=\, \int_{\mathbb{R}^d} \big(f * \widetilde{g}\big)(v)\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(v) \,=\, \mu^{(2)}_{\mathrm{red}}\big(f * \widetilde{g}\big),
\end{aligned} \tag{6.6}$$

so that

$$\begin{aligned}
\mu^{(2)}_{\mathrm{red}}\big(f * \widetilde{f}\,\big) \,&=\, \mu^{(2)}\big(f \otimes \overline{f}\big) \,=\, \mathbb{E}\Big[\textstyle\int f\,\mathrm{d}\varPhi \int \overline{f}\,\mathrm{d}\overline{\varPhi}\Big]\\
&=\, \mathbb{E}\left[\int f\,\mathrm{d}\varPhi\,\overline{\int f\,\mathrm{d}\varPhi}\right] \,=\, \mathbb{E}\left[|\varPhi(f)|^2\right] \geq 0.
\end{aligned}$$

Remark 6.1 (see also [4, Rem. 13]**).** One can alternatively define

$$\mu^{(2,\mathrm{alt})}(A \times A') \,=\, \mathbb{E}\big[\,\overline{\varPhi(A)}\varPhi(A')\big] \qquad \text{for bounded } \; A, A' \in \mathcal{B}(\mathbb{R}^d),$$

and then obtain $\mu_{\mathrm{red}}^{(2,\mathrm{alt})}$ from this as above, via

$$\begin{aligned}\int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2,\mathrm{alt})}(x,y) &= \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u+v,u)\,\mathrm{d}\mu_{\mathrm{red}}^{(2,\mathrm{alt})}(v)\,\mathrm{d}\lambda(u)\\ &= \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u,u-v)\,\mathrm{d}\mu_{\mathrm{red}}^{(2,\mathrm{alt})}(v)\,\mathrm{d}\lambda(u).\end{aligned}$$

Then, we have $\mu^{(2,\mathrm{alt})} = \overline{\mu^{(2)}}$ and $\mu_{\mathrm{red}}^{(2,\mathrm{alt})} = \overline{\mu_{\mathrm{red}}^{(2)}}$. Since

$$\begin{aligned}\int_{\mathbb{R}^d\times\mathbb{R}^d} f(y,x)\,\mathrm{d}\mu^{(2)}(x,y) &= \int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\overline{\mu^{(2)}}(x,y)\\ &= \int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2,\mathrm{alt})}(x,y),\end{aligned}$$

we see that the alternative choice of factoring out the shift invariance in Eq. (6.4), namely integrating $f(u,u+v)$ on the right-hand side of this equation, leads to $\mu_{\mathrm{red}}^{(2,\mathrm{alt})}$, where

$$\int_{\mathbb{R}^d\times\mathbb{R}^d} f(x,y)\,\mathrm{d}\mu^{(2)}(x,y) = \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} f(u,u+v)\,\mathrm{d}\mu_{\mathrm{red}}^{(2,\mathrm{alt})}(v)\,\mathrm{d}\lambda(u)\,. \tag{6.7}$$

We choose the definitions as in Eqs. (6.3) and (6.4) because these fit well to the formulation of the limit in Eq. (6.8) below. Note that, in the real-valued case, $\mu_{\mathrm{red}}^{(2)}$ and $\mu_{\mathrm{red}}^{(2,\mathrm{alt})}$ agree.

The 'complex-valued' analogue of [4, Thm. 5] now reads as follows.

Theorem 6.2. *Assume that conditions* (6.1) *and* (6.2) *are satisfied, and let* $\varPhi_n := \varPhi|_{B_n}$ *denote the restriction of* $\varPhi$ *to the open ball of radius* n *around* 0*. Then, the natural autocorrelation of* $\varPhi$*, which is defined with an averaging sequence of nested, centred balls, almost surely exists and satisfies*

$$\gamma^{(\varPhi)} := \lim_{n\to\infty} \frac{\varPhi_n * \widetilde{\varPhi_n}}{\lambda(B_n)} = \lim_{n\to\infty} \frac{\varPhi_n * \widetilde{\varPhi}}{\lambda(B_n)} = \mu_{\mathrm{red}}^{(2)}, \tag{6.8}$$

where the limit refers to the vague topology. In particular, the autocorrelation is non-random.

Proof. The proof is a suitable 'complex-valued interpretation' of the proof of [4, Thm. 5]. Fix a continuous function $f\colon \mathbb{R}^d \longrightarrow \mathbb{C}$ with compact support. We have to check that

$$\frac{1}{\lambda(B_n)}\bigl(\varPhi_n * \widetilde{\varPhi_n}\bigr)(f) \xrightarrow{\;n\to\infty\;} \mu_{\mathrm{red}}^{(2)}(f) \qquad \text{(a.s.)}. \tag{6.9}$$

Since both sides are locally finite (complex-valued) measures, it actually suffices to check Eq. (6.9) for real-valued f. For $x\in\mathbb{R}^d$, define

$$F(x) := \int_{\mathbb{R}^d} f(x-y)\,\mathrm{d}\overline{\varPhi}(y) = \int_{\mathbb{R}^d} f(x+y)\,\mathrm{d}\widetilde{\varPhi}(y).$$

Clearly, F inherits stationarity and ergodicity from $\varPhi$, wherefore F is a (complex-valued) ergodic random function on $\mathbb{R}^d$ in the sense that shift-invariant events for F have 'trivial' probabilities (0 or 1), and we obtain

$$\mathbb{E}\left[\int_A \bigl|F(x)\bigr| \,\mathrm{d}|\varPhi|(x)\right] \,<\, \infty$$

for any bounded and measurable $A \subset \mathbb{R}^d$.

Define a (complex-valued) additive covariant spatial process $X(A)$ in the sense of [66], indexed by a bounded and measurable $A \subset \mathbb{R}^d$, via

$$X(A) \,:=\, \int_A F(x)\,\mathrm{d}\varPhi(x)\,.$$

Covariant in this context means that X behaves 'naturally' under translations: When $\mathbb{R}^d$ acts on X via $(T_u X)(A) := \int_A F(x)\,\mathrm{d}(T_u\varPhi)(x)$, for $u \in \mathbb{R}^d$, then $(T_u X)(A+u) = X(A)$.

Decomposing X into its real and imaginary parts (by decomposing F and $\varPhi$ and suitably grouping terms) we can apply [66, Cor. 4.9] to obtain a.s.

$$\begin{aligned}
\lim_{n\to\infty} \frac{1}{\lambda(B_n)} \bigl(\varPhi_n * \widetilde{\varPhi}\bigr)(f) \,&=\, \lim_{n\to\infty} \frac{1}{\lambda(B_n)} \int_{B_n} F(x)\,\mathrm{d}\varPhi(x) \,=\, \lim_{n\to\infty} \frac{X(B_n)}{\lambda(B_n)} \\
&=\, \mathbb{E}\left[\frac{X(B_1)}{\lambda(B_1)}\right] \,=\, \frac{1}{\lambda(B_1)}\,\mathbb{E}\left[\int_{B_1}\int_{\mathbb{R}^d} f(x-y)\,\mathrm{d}\overline{\varPhi}(y)\,\mathrm{d}\varPhi(x)\right] \\
&=\, \frac{1}{\lambda(B_1)} \int_{\mathbb{R}^d\times\mathbb{R}^d} \mathbf{1}_{B_1}(x)\, f(x-y)\,\mathrm{d}\mu^{(2)}(x,y) \\
&=\, \frac{1}{\lambda(B_1)} \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \mathbf{1}_{B_1}(x)\, f(z)\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}(z)\,\mathrm{d}\lambda(x) \,=\, \int_{\mathbb{R}^d} f\,\mathrm{d}\mu^{(2)}_{\mathrm{red}}.
\end{aligned}$$

The difference between $\varPhi_n * \widetilde{\varPhi}$ and $\varPhi_n * \widetilde{\varPhi_n}$ is a (random) 'boundary term' that almost surely vanishes in the limit as $n \to \infty$. To prove this formally, decompose $\varPhi = \varPhi_{\Re} + \mathrm{i}\varPhi_{\Im}$, $\widetilde{\varPhi} = \widetilde{\varPhi_{\Re}} - \mathrm{i}\widetilde{\varPhi_{\Im}}$ and then argue as in the proof of [4, Thm. 3] for each of the four terms appearing in $\varPhi_n * \bigl(\widetilde{\varPhi} - \widetilde{\varPhi_n}\bigr)$. □

Remark 6.3. Theorem 6.2 allows to reformulate Theorem 4 and Corollary 1 from [4] for complex-valued clusters as follows. If $\varPhi$ is a stationary ergodic point process, i.e., $\varPhi$ is a random sum of Dirac measures, with distribution P satisfying Eq. (6.2), and if we replace each point independently by a random complex-valued measure with distribution Q, then the formulas describing the autocorrelation and the diffraction of the resulting cluster process given in [4, Thm. 4 and Cor. 1] continue to hold.

Let us also mention that, by specialising $\varPhi$ to a renewal process, Theorem 6.2 allows to recover Eq. (4.1) and, in particular, Theorem 4.1 from this more general perspective; see [4] for further details, and how this can be used to formulate the renewal process also for more general 'dropping' distributions.

6.1. A 'Palm-type distribution' for complex-valued random measures

In the case of a positive random measure $\varPhi$, Eq. (6.8) can be interpreted via the Palm distribution P_0 of the law of $\varPhi$, which is a probability measure on locally finite measures (intuitively, the law of $\varPhi$ viewed relative to a typical point of its support) via

$$\mu_{\mathrm{red}}^{(2)} \,=\, \rho\, I_{P_0} \tag{6.10}$$

where $\rho > 0$ is the intensity and I_{P_0} the first moment measure of P_0; compare [4, Sec. 5.2]. This interpretation breaks down in general in the signed or complex-valued case because $\mu_{\mathrm{red}}^{(2)}$ will not be a positive measure. One way to extend this line of thought is to re-interpret the Palm distributions in a way suited for complex-valued random measures as follows.

Recalling the structure of the polar decomposition, the random measure $\varPhi$ can equivalently be described via $(|\varPhi|, \varPhi_{\mathrm{ph}})$, where $|\varPhi|$ is the total variation measure and $\varPhi_{\mathrm{ph}}\colon \mathbb{R}^d \longrightarrow [0, 2\pi)$ the 'phase function':

$$\int_{\mathbb{R}^d} f(x)\, \mathrm{d}\varPhi(x) \,=\, \int_{\mathbb{R}^d} f(x)\, \mathrm{e}^{\mathrm{i}\varPhi_{\mathrm{ph}}(x)}\, \mathrm{d}|\varPhi|(x). \tag{6.11}$$

Note that $\varPhi \mapsto (|\varPhi|, \varPhi_{\mathrm{ph}})$ is measurable, so $(|\varPhi|, \varPhi_{\mathrm{ph}})$ is in fact a random variable. Define a positive σ-finite measure $\mathcal{C}$ on $\mathbb{R}^d \times \mathcal{M}$ (this is the equivalent of the so-called Campbell measure for the complex-valued context and agrees with the usual Campbell measure if $\varPhi$ is a positive random measure) via

$$\int_{\mathbb{R}^d \times \mathcal{M}} g(x, \varphi)\, \mathrm{d}\mathcal{C}(x, \varphi) \,:=\, \mathbb{E}\Big[\int_{\mathbb{R}^d} g\big(x, \mathrm{e}^{-\mathrm{i}\varPhi_{\mathrm{ph}}(x)}\varPhi\big)\, \mathrm{d}|\varPhi|(x) \Big],$$

whenever the right-hand side is defined (which will for instance always be the case when g is measurable and non-negative). By the shift invariance of $\varPhi$, and hence that of $|\varPhi|$, the projection of $\mathcal{C}$ to $\mathbb{R}^d$ is ρ times Lebesgue measure (with $\rho \in [0, \infty)$ being the intensity of $|\varPhi|$), hence there is a family of probability measures P_x on $\mathcal{M}$, with $P_x \in \mathcal{P}(\mathcal{M})$ for all $x \in \mathbb{R}^d$, so that we can disintegrate (compare [48, Thm. 15.3.3])

$$\int g\, \mathrm{d}\mathcal{C} \,=\, \int_{\mathbb{R}^d} \int_{\mathcal{P}(\mathcal{M})} g(x, \varphi)\; \mathrm{d}P_x(\varphi)\, \rho\, \mathrm{d}\lambda(x). \tag{6.12}$$

Definition 6.4. We call the elements of the family $\big\{P_x \mid x \in \mathbb{R}^d\big\}$ the *Palm distributions* in the complex-valued case.

Let, for $A \subset \mathbb{R}^d$ bounded and measurable,

$$I_{P_x}(A) \,:=\, \int_{\mathcal{M}} \varphi(A)\, \mathrm{d}P_x(\varphi)$$

be the expectation (or first moment) measure of P_x. By shift invariance, we have $P_x = T_x P_0$, $x \in \mathbb{R}^d$, and hence $I_{P_x} = T_x I_{P_0}$. The connection between the (reduced) second moment measure and the Palm distribution carries over to the complex-valued case as follows.

Proposition 6.5. *For the extended definition of the Palm distribution, one has*

$$\mu^{(2)}_{\mathrm{red}} \,=\, \rho\, I_{P_0}\,,$$

so Eq. (6.10) *also holds in this case.*

Sketch of proof. Consider $g(x,\varphi) = \mathbf{1}_{A'}(x)\varphi(A)$ with $\mathbf{1}$ denoting the characteristic function and with $A, A' \subset \mathbb{R}^d$ bounded and measurable. Then,

$$\begin{aligned}\int g\,\mathrm{d}\mathcal{C} \,&=\, \mathbb{E}\Big[\int_{\mathbb{R}^d} \mathbf{1}_{A'}(x)\mathrm{e}^{-\mathrm{i}\varPhi_{\mathrm{ph}}(x)}\varPhi(A)\,\mathrm{d}|\varPhi|(x)\Big] \qquad (6.13)\\ &=\, \mathbb{E}\Big[\varPhi(A)\int_{\mathbb{R}^d} \mathbf{1}_{A'}(x)\mathrm{e}^{-\mathrm{i}\varPhi_{\mathrm{ph}}(x)}\,\mathrm{d}|\varPhi|(x)\Big] \,=\, \mathbb{E}\big[\varPhi(A)\overline{\varPhi(A')}\,\big]\end{aligned}$$

by definition, whereas the disintegration formula yields

$$\begin{aligned}\int g\,\mathrm{d}\mathcal{C} \,&=\, \int_{\mathbb{R}^d}\int_{\mathcal{M}} \mathbf{1}_{A'}(x)\varphi(A)\;\mathrm{d}P_x(\varphi)\,\rho\,\mathrm{d}\lambda(x) \,=\, \rho\int_{A'} I_{P_x}(A)\,\mathrm{d}\lambda(x)\\ &=\, \rho\int_{A'} I_{P_0}(A-x)\,\mathrm{d}\lambda(x) \,=\, \rho\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \mathbf{1}_{A'}(x)\mathbf{1}_{A-x}(y)\,\mathrm{d}I_{P_0}(y)\,\mathrm{d}\lambda(x)\\ &=\, \rho\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \mathbf{1}_{A'}(x)\mathbf{1}_{A}(y+x)\,\mathrm{d}\lambda(x)\,\mathrm{d}I_{P_0}(y) \qquad (6.14)\\ &=\, \rho\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \mathbf{1}_{A'}(-x)\mathbf{1}_{A}(y-x)\,\mathrm{d}\lambda(x)\,\mathrm{d}I_{P_0}(y) \,=\, \rho\, I_{P_0}\big(\mathbf{1}_A * \widetilde{\mathbf{1}_{A'}}\big).\end{aligned}$$

Comparing Eqs. (6.13)–(6.14) with Eq. (6.6) yields the claim. □

If $\varPhi$ is ergodic, the viewpoint that P_0 describes the configuration relative to a point in the support drawn according to $\varPhi$ is corroborated by

$$\frac{1}{\lambda(B_n)}\int_{B_n} g\big(\mathrm{e}^{-\mathrm{i}\varPhi_{\mathrm{ph}}(x)}T_{-x}\varPhi\big)\,\mathrm{d}|\varPhi|(x) \;\xrightarrow{\;n\to\infty\;}\; \int_{\mathcal{M}} g(\varphi)\,\mathrm{d}P_0(\varphi) \quad \text{(a.s.)}$$

for any bounded measurable $g\colon \mathcal{M} \longrightarrow \mathbb{R}$.

The viewpoint of (possibly complex-valued) ergodic random measures for diffraction is a useful one since it provides a connection to the large literature on random measures and on stochastic geometry; see [30, 49, 51, 52, 62, 4] and references therein, as well as [75] for a recent generalisation that can also be considered from the diffraction point of view. However, our approach also shows a limitation that one encounters when trying to infer properties of a random configuration of scatterers from its kinematic diffraction: As is evident from Eq. (6.8) in Theorem 6.2, the only 'datum' from a random $\varPhi$ visible in its autocorrelation, and hence also in the corresponding diffraction, is the second moment measure. It is well known that second moments are generally insufficient to determine the distribution of $\varPhi$ unless further structural properties are known. This inverse problem is known as the homometry problem in crystallography and the inference problem in the theory of stochastic processes.

7. Outlook

Our exposition provides a snapshot of the present knowledge about systems with continuous diffraction components; see [12, Chs. 10 and 11] as well as [4, 15] for additional examples, and [20, 15] for connections with the dynamical spectrum. Nevertheless, as is apparent from a comparison with the pure point diffraction case [23, 12, 58, 77, 76], the status of general results is lagging behind. Even for many important examples, some of the most obvious questions are still open from a mathematical point of view. In particular, this is so for random tiling ensembles in dimensions $d \geq 2$, or for equilibrium systems just beyond the complexity of the (planar) Ising model.

Apart from the systems considered here, an interesting class is provided by random substitution and inflation systems, as introduced in [41]. The randomness present here is compatible with the long-range order of Meyer sets with entropy [21, 64, 65], which means that one obtains interesting mixtures of pure point and absolutely continuous diffraction measures. Though this direction has not attracted much attention so far, it is both tractable and practically relevant.

From a more general perspective, one lacks some kind of analogue to the key theorems in pure point diffraction (such as the Poisson summation formula or the Halmos–von Neumann theorem). While there is at least the theory of Riesz products [80, 68] for self-similar systems with singular spectra, a general approach to stochastic systems is only at its beginning. Methods from point process theory [30], such as the Palm measure and its connection to the autocorrelation (via its intensity measure), look promising, but have not produced many concrete results so far. The latter, however, are needed to make some progress with the complicated inverse problem for such systems. Though there is substantial knowledge from the inference approach [49], it is not clear at present how this can be used, and how reasonable restrictions could be included.

Acknowledgment

It is our pleasure to thank Aernout van Enter, Holger Kösters, Daniel Lenz, Robert Moody and Tom Ward for interesting discussions and for their cooperation on some of the papers that form the basis for this review. This work was supported by the German Research Council (DFG), within the CRC 701.

References

[1] J.-P. Allouche and J. Shallit, *Automatic Sequences: Theory, Applications, Generalizations*. Cambridge University Press, Cambridge, 2003.

[2] G.W. Anderson, A. Guionnet and O. Zeitouni, *An Introduction to Random Matrices*. Cambridge University Press, Cambridge, 2010.

[3] M. Baake, *Diffraction of weighted lattice subsets*, Can. Math. Bulletin **45** (2002), 483–498; arXiv:math.MG/0106111.

[4] M. Baake, M. Birkner and R.V. Moody, *Diffraction of stochastic point sets: Explicitly computable examples*, Commun. Math. Phys. **293** (2010), 611–660; arXiv:0803.1266.

[5] M. Baake, F. Gähler and U. Grimm, *Spectral and topological properties of a family of generalised Thue–Morse sequences*, J. Math. Phys. **53** (2012), 032701 (24pp); arXiv:1201.1423.

[6] M. Baake, F. Götze, C. Huck and T. Jakobi, *Radial spacing distributions from planar points sets*, Acta Crystallogr. A **70** (2014), 472–482; arXiv:1402.2818.

[7] M. Baake and U. Grimm, *The singular continuous diffraction measure of the Thue–Morse chain*, J. Phys. A.: Math. Theor. **41** (2008), 422001 (6 pp.); arXiv:0809.0580.

[8] M. Baake and U. Grimm, *Kinematic diffraction is insufficient to distinguish order from disorder*, Phys. Rev. B **79** (2009), 020203(R) (4 pp.) and **80**, 029903(E) (1 p.); arXiv:0810.5750.

[9] M. Baake and U. Grimm, *Kinematic diffraction from a mathematical viewpoint*, Z. Krist. **226** (2011), 711–725; arXiv:1105.0095.

[10] M. Baake and U. Grimm, *A comment on the relation between diffraction and entropy*, Entropy **14** (2012), 856–864; arXiv:1205.0392.

[11] M. Baake and U. Grimm, *Squirals and beyond: Substitution tilings with singular continuous spectrum*, Ergod. Th. & Dynam. Syst. **34** (2014), 1077–1102; arXiv:1205.1384.

[12] M. Baake and U. Grimm, *Aperiodic Order. Volume 1: A Mathematical Invitation*, Cambridge University Press, Cambridge, 2013.

[13] M. Baake and M. Höffe, *Diffraction of random tilings: Some rigorous results*, J. Stat. Phys. **99** (2000), 219–261; arXiv:math-ph/9904005.

[14] M. Baake and H. Kösters, *Random point sets and their diffraction*, Philos. Mag. **91** (2011), 2671–2679; arXiv:1007.3084.

[15] M. Baake, H. Kösters and R.V. Moody, *Diffraction theory of point processes: Systems with clumping and repulsion*, J. Stat. Phys., in press; DOI:10.1007/s10955-014-1178-5; arXiv:1405.4255.

[16] M. Baake and D. Lenz, *Dynamical systems on translation bounded measures: Pure point dynamical and diffraction spectra*, Ergodic Th. & Dynam. Syst. **24** (2004), 1867–1893; arXiv:math.DS/0302231.

[17] M. Baake and D. Lenz, *Deformation of Delone dynamical systems and topological conjugacy*, J. Fourier Anal. Appl. **11** (2005), 125–150; arXiv:math.DS/0404155.

[18] M. Baake, D. Lenz and R.V. Moody, *Characterization of model sets by dynamical systems*, Ergodic Th. & Dynam. Syst. **27** (2007), 341–382; arXiv:math.DS/0511648.

[19] M. Baake, D. Lenz and C. Richard, *Pure point diffraction implies zero entropy for Delone sets with uniform cluster frequencies*, Lett. Math. Phys. **82** (2007), 61–77; arXiv:0706.1677.

[20] M. Baake, D. Lenz and A.C.D. van Enter, *Dynamical versus diffraction spectrum for structures with finite local complexity*, Ergod. Th. & Dynam. Syst., in press; DOI:10.1017/etds.2014.28; arXiv:1307.7518.

[21] M. Baake and M. Moll, *Random noble means substitutions*, in: *Aperiodic Crystals*, eds. S. Schmid, R.L. Withers and R. Lifshitz (Springer, Dordrecht, 2013), pp. 19–27; arXiv:1210.3462.

[22] M. Baake and R.V. Moody, *Diffractive point sets with entropy*, J. Phys. A: Math. Gen. **31** (1998), 9023–9039; arXiv:math-ph/9809002.

[23] M. Baake and R.V. Moody, *Weighted Dirac combs with pure point diffraction*, J. reine angew. Math. (Crelle) **573** (2004), 61–94; arXiv:math.MG/0203030.

[24] M. Baake, R.V. Moody and M. Schlottmann, *Limit-(quasi)periodic point sets as quasicrystals with p-adic internal spaces*, J. Phys. A: Math. Gen. **31** (1998), 5755–5765; arXiv:math-ph/9901008.

[25] M. Baake and A.C.D. van Enter, *Close-packed dimers on the line: Diffraction versus dynamical spectrum*, J. Stat. Phys. **143** (2011), 88–101; arXiv:1011.1628.

[26] M. Baake and T. Ward, *Planar dynamical systems with pure Lebesgue diffraction spectrum*, J. Stat. Phys. **140** (2010), 90–102; arXiv:1003.1536.

[27] C. Berg and G. Forst, *Potential Theory on Locally Compact Abelian Groups*. Springer, Berlin, 1975.

[28] I.P. Cornfeld, S.V. Fomin and Ya.G. Sinai, *Ergodic Theory*, Springer, New York, 1982.

[29] J.M. Cowley, *Diffraction Physics*, 3rd ed., North-Holland, Amsterdam, 1995.

[30] D.J. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, Springer, New York, 1988.

[31] D. Damanik, M. Embree and A. Gorodetski, *Spectral properties of Schrödinger operators arising in the study of quasicrystals*, in: Mathematics of Aperiodic Order, J. Kellendonk, D. Lenz and J. Savinien (eds.), Progress in Mathematics vol. 309, Springer Basel AG, ISBN 978-3-0348-0902-3, 307–370; arXiv:1210.5753.

[32] X. Deng and R.V. Moody, *How model sets can be determined by their two-point and three-point correlations*, J. Stat. Phys. **135** (2009), 621–637; arXiv:0901.4381.

[33] J. Dieudonné, *Treatise on Analysis*, Vol. II, Academic Press, New York, 1970.

[34] F. Dyson, *Statistical theory of the energy levels of complex systems. III*, J. Math. Phys. **3** (1962), 166–175.

[35] M. Einsiedler and T. Ward, *Ergodic Theory, with a View towards Number Theory*, Springer, London, 2011.

[36] N. Etemadi, *An elementary proof of the strong law of large numbers*, Z. Wahrscheinlichkeitsth. verw. Geb. **55** (1981), 119–122.

[37] W. Feller, *An Introduction to Probability Theory and Its Applications*, vol. II, 2nd ed., Wiley, New York, 1971.

[38] P.J. Forrester, *Log-Gases and Random Matrices*, Princeton University Press, Princeton, 2010.

[39] N.P. Frank, *Substitution sequences in $\mathbb{Z}^d$ with a non-simple Lebesgue component in the spectrum*, Ergod. Th. & Dynam. Syst. **23** (2003), 519–532.

[40] N.P. Frank, *Multi-dimensional constant-length substitution sequences*, Topol. Appl. **152** (2005), 44–69.

[41] C. Godrèche and J.-M. Luck, *Quasiperiodicity and randomness in tilings of the plane*, J. Stat. Phys. **55** (1989), 1–28.

[42] J.-B. Goueré, *Diffraction and Palm measure of point processes*, Comptes Rendus Acad. Sci. (Paris) **342** (2003), 141–146; arXiv:math.PR/0208064.

[43] C. Henley, *Random tiling models*, in: *Quasicrystals: The State of the Art*, 2nd ed., eds. D.P. DiVincenzo and P.J. Steinhardt (World Scientific, Singapore, 1999), pp. 459–560.

[44] A. Hof, *On diffraction by aperiodic structures*, Commun. Math. Phys. **169** (1995), 25–43.

[45] M. Höffe, *Diffraction of the dart-rhombus random tiling*, Math. Sci. Eng. A **294–296** (2000), 373–376; arXiv:math-ph/9911014.

[46] M. Höffe, *Diffraktionstheorie stochastischer Parkettierungen*, Shaker, Aachen, 2001.

[47] S. Kakutani, *Strictly ergodic symbolic dynamical systems*, in: *Proc. 6th Berkeley Symposium on Math. Statistics and Probability*, eds. L.M. LeCam, J. Neyman and E.L. Scott (Univ. of California Press, Berkeley, 1972), pp. 319–326.

[48] O. Kallenberg, *Random Measures*, 3rd ed., Akademie-Verlag, Berlin, 1983.

[49] A.F. Karr, *Point Processes and their Statistical Inference*, 2nd ed., Dekker, New York, 1991.

[50] Y. Katznelson, *An Introduction to Harmonic Analysis*, 3rd ed., Cambridge University Press, New York, 2004.

[51] C. Külske, *Universal bounds on the selfaveraging of random diffraction measures*, Probab. Th. Relat. Fields **126** (2003), 29–50; arXiv:math-ph/0109005.

[52] C. Külske, *Concentration inequalities for functions of Gibbs fields with application to diffraction and random Gibbs measures*, Commun. Math. Phys. **239** (2003), 29–51.

[53] S. Lang, *Real and Functional Analysis*, 3rd ed., Springer, New York, 1993.

[54] F. Ledrappier, *Un champ markovien peut être d'entropie nulle et mélangeant*, C. R. Acad. Sci. Paris Sér. A-B **287** (1987), A561–A563.

[55] J.-Y. Lee, R.V. Moody and B. Solomyak, *Pure point dynamical and diffraction spectra*, Ann. H. Poincaré **3** (2002), 1003–1018; arXiv:0910.4809.

[56] D. Lenz, *Continuity of eigenfunctions of uniquely ergodic dynamical systems and intensity of Bragg peaks*, Commun. Math. Phys. **287** (2009), 225–258; arXiv:math-ph/0608026.

[57] D. Lenz and R.V. Moody, *Extinctions and correlations for uniformly discrete point processes with pure point dynamical spectra*, Commun. Math. Phys. **289** (2009), 907–923; arXiv:0902.0567.

[58] D. Lenz and R.V. Moody, *Stationary processes with pure point diffraction*, preprint; arXiv:1111.3617.

[59] D. Lenz and C. Richard, *Pure point diffraction and cut and project schemes for measures: The smooth case*, Math. Z. **256** (2007), 347–378; arXiv:math.DS/0603453.

[60] D. Lenz and N. Strungaru, *Pure point spectrum for measure dynamical systems on locally compact Abelian groups*, J. Math. Pures Appl. **92** (2009), 323–341; arXiv:0704.2498.

[61] K. Mahler, *The spectrum of an array and its application to the study of the translation properties of a simple class of arithmetical functions. Part II: On the translation properties of a simple class of arithmetical functions*, J. Math. Massachusetts **6** (1927), 158–163.

[62] K. Matzutt, *Diffraction of Point Sets with Structural Disorder*, PhD thesis, Bielefeld University, Bielefeld, 2010.

[63] M.L. Mehta, *Random Matrices*, 3rd ed., Elsevier, Amsterdam, 2004.

[64] M. Moll, *On a Family of Random Noble Means Substitutions*, PhD thesis, Bielefeld University, Bielefeld, 2013.

[65] M. Moll, *Diffraction of random noble means words*, J. Stat. Phys. **156** (2014), 1221–1236; arXiv:1404.7411.

[66] X.X. Nguyen and H. Zessin, *Ergodic theorems for spatial processes*, Z. Wahrsch. verw. Gebiete **48** (1979), 133–158.

[67] N. Pytheas Fogg, *Substitutions in Dynamics, Arithmetics and Combinatorics*, LNM 1794, Springer, Berlin, 2002.

[68] M. Queffélec, *Substitution Dynamical Systems – Spectral Analysis*. 2nd ed., LNM 1294, Springer, Berlin, 2010.

[69] C. Richard, *An alternative view on random tilings*, J. Phys. A: Math. Gen. **32** (1999), 8823–8829; arXiv:cond-mat/9907262.

[70] C. Richard, *Statistische Physik stochastischer Parkettierungen*, UFO Verlag, Allensbach, 1999.

[71] C. Richard, *Dense Dirac combs in Euclidean space with pure point diffraction*, J. Math. Phys. **44** (2003), 4436–4449; arXiv:math-ph/0302049.

[72] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, 2nd ed., Academic Press, San Diego, 1980.

[73] K. Schmidt, *Dynamical Systems of Algebraic Origin*, Birkhäuser, Basel, 1995.

[74] J. Slawny, *Ergodic properties of equilibrium states*, Commun. Math. Phys. **80** (1981), 477–483.

[75] J. Teichmann, F. Ballani and K.G. van den Boogaart, *Generalizations of Matérn's hard-core point processes*, Spatial Statistics **3** (2013), 33–53; arXiv:1209.2566.

[76] V. Terauds, *The inverse problem of pure point diffraction – examples and open questions*, J. Stat. Phys. **152** (2013), 954–968; arXiv:1303.3260.

[77] V. Terauds and M. Baake, *Some comments on the inverse problem of pure point diffraction*, in: *Aperiodic Crystals*, eds. S. Schmid, R.L. Withers and R. Lifshitz (Springer, Dordrecht, 2013), pp. 35–41; arXiv:1210.3460.

[78] A.C.D. van Enter and J. Miękisz, *How should one define a (weak) crystal?* J. Stat. Phys. **66** (1992), 1147–1153.

[79] N. Wiener, *The spectrum of an array and its application to the study of the translation properties of a simple class of arithmetical functions. Part I: The spectrum of an array*, J. Math. Massachusetts **6** (1927), 145–157.

[80] Zygmund A, *Trigonometric Series*, 3rd ed., Cambridge University Press, Cambridge, 2002.

Michael Baake
Fakultät für Mathematik
Universität Bielefeld
Postfach 100131
D-33501 Bielefeld, Germany
e-mail: mbaake@math.uni-bielefeld.de

Matthias Birkner
Institut für Mathematik
Johannes-Gutenberg-Universität
Staudingerweg 9
D-55099 Mainz, Germany
e-mail: birkner@mathematik.uni-mainz.de

Uwe Grimm
Department of Mathematics and Statistics
The Open University
Walton Hall
Milton Keynes MK7 6AA, UK
e-mail: uwe.grimm@open.ac.uk

Progress in Mathematics, Vol. 309, 33–72
© 2015 Springer Basel

On the Pisot Substitution Conjecture

S. Akiyama, M. Barge, V. Berthé, J.-Y. Lee and A. Siegel

Abstract. Our goal is to present a unified and reasonably complete account of the various conjectures, known as Pisot conjectures, that assert that certain dynamical systems arising from substitutions should have pure discrete dynamical spectrum. We describe the various contexts (symbolic, geometrical, arithmetical) in which substitution dynamical systems arise and review the relevant properties of these systems. The Pisot Substitution Conjecture is stated in each context and the relationships between these statements, and with several related conjectures, are discussed. We survey the special cases in which the Pisot Substitution Conjecture has been verified and present algorithmic procedures for checking pure discrete spectrum. We conclude with a discussion of possible extensions to higher dimensions.

Mathematics Subject Classification (2010). Primary 37B50; Secondary 37B10, 37B05.

Keywords. Tiling space, substitution, Pisot number, discrete spectrum, cut-and-project scheme, beta-numeration, Delone set, Meyer set, model set, coincidence.

1. Introduction

Substitutive dynamics. Substitutions are replacement rules that can be of either a symbolic or a geometric nature: by iteration, they produce hierarchically ordered structures (infinite words, point sets, tilings) that display strong self-similarity properties. A symbolic substitution is a morphism on a free monoid defined by replacing letters by finite words, while a geometric substitution inflates each tile in some finite collection of tiles and subdivides the inflated tile into translates of the original tiles (the tiles are like letters and the inflated and subdivided tiles are like higher-dimensional words). Iterating a substitution on a letter (or tile) produces

This work was supported by the Agence Nationale de la Recherche through contract ANR-13-BS02-0003 and ANR-12-IS01-0002-01. Also this work was supported by the National Research Foundation of Korea (NRF) Grant funded by the Korean Government(MSIP)(2014004168) and Korea Institute for Advanced Study (KIAS).

longer and longer words (or larger and larger 'patches'), any subword (or subpatch) of such a word (or patch) is said to be *allowed* for the substitution. The space of all bi-infinite words (or tilings), all of whose finite subwords (or patches) are allowed for a substitution is called a *substitutive system* (or *substitution tiling space*).

The idea of a substitution as a replacement rule is central in symbolic dynamics and tiling theory (see, for instance, [96, 95, 99, 113, 16]). Indeed, substitutions are closely related to the process of induction. For a measure-preserving dynamical system, the Poincaré recurrence theorem guarantees that almost every point in any given set eventually returns to that set. An *induced* system is then defined by first return to the set and a system is *self-induced* if there is a subset with positive measure for which the induced system is isomorphic to the original system. Self-induced dynamics underlies periodic expansions with respect to various algorithms. For example, the continued fraction expansion of a quadratic irrational is derived from successive induction on irrational rotation of the circle. In this case the maximal eigenvalue of the matrix associated with the inverse application for the self-inducing structure is a quadratic Pisot unit and this fact is used to effectively find the fundamental unit of a real quadratic number field. Similar phenomena occur for piecewise isomorphisms and outer billiards.

The field of symbolic dynamics has its origins in the coding of concrete (geometrical) dynamical systems (see, *e.g.*, [90, 91]). Given a partition of the phase space of a dynamical system there is an associated space consisting of the collection of all itineraries of the system with respect to the partition. The symbolic dynamical system consisting of the shift map on the space of itineraries is then a model of the original system. We are interested here in the inverse problem: given a symbolic dynamical system or a tiling system, is there a geometrical dynamical system that it codes? More particularly, we are interested in the geometrical interpretation of symbolic systems consisting of infinite words (or tilings) created by a substitution. In order to have a chance to find the nicest sort of geometrical interpretation – translation on a compact abelian group – of a substitutive symbolic system, it is necessary to restrict to substitutions that create a hierarchical structure with a significant amount of long range order. As translation on a compact abelian group is almost periodic, the substitutions we consider must create sequences, or tilings, with a similar structure.

A (primitive) substitution ϕ stretches words, on average, by some factor λ: if w is a long (allowed) word then the word $\phi(w)$ has roughly λ times the number of letters as has w. The number λ is the *expansion factor* of ϕ. For the infinite words in the substitutive system associated with ϕ to be shift-periodic, λ must be an integer (*e.g.*, $\phi(a) = ab$, $\phi(b) = ab$, $\lambda = 2$). It turns out that for the infinite words making up the substitutive system associated with ϕ to be nearly enough periodic in order that the substitutive system be a coding of a translation on a compact abelian group, it is necessary that higher and higher powers of λ are more and more nearly integers. Pisot characterized such λ ([92]): if $\lambda > 1$ is an algebraic integer, then the distance from λ^n to the nearest integer goes to zero as n goes to infinity if and only if all of the algebraic conjugates of λ (other than λ)

lie strictly inside the unit circle. Such λ are called *Pisot numbers* (or sometimes *Pisot–Vijayaraghavan numbers*).

Pisot substitutions. For the above reasons, we restrict ourselves here to substitutions whose expansion factor is a Pisot number. Tiling substitutions also have an associated linear expansion: for one-dimensional tiling substitutions we will assume that this expansion is a Pisot number and in higher dimensions that it has the *Pisot property* (see Section 7.1 for a definition).

By a geometrical interpretation of a symbolic (or tiling) system we mean a factor map from the system onto some system of a geometrical nature. If the geometrical system is translation on a compact abelian group, the existence of such a (non-trivial) factorization is equivalent to the symbolic system having a non-trivial dynamical spectrum. A dynamical system has *pure discrete spectrum* if it factors almost everywhere one-to-one onto a translation on a compact abelian group. The connections between Pisot substitutions and discrete spectrum first appear in [42, 43, 112, 97, 117] and it is shown in [113] that for a one-dimensional substitutive system to have pure discrete spectrum it is necessary that the expansion λ be a Pisot number. Our main focus in this chapter is on the question: what conditions must be placed on a Pisot substitution in order to guarantee that the associated substitutive system (or tiling system) has pure discrete spectrum? The Pisot Substitution Conjecture, and its variants, are proposed answers.

Organization of the chapter. Section 2 introduces various types of substitutions and their associated dynamical systems. These include substitutions acting on words, tilings and point sets. We also consider the related framework of beta-numeration.

In Section 3 we discuss the notions of discreteness of the dynamical and diffraction spectra of the substitutive systems. In particular, the Meyer property is introduced in Section 3.2, and in Section 3.3 we give a general definition of cut-and-project schemes and model sets.

The various (one-dimensional) Pisot Substitution Conjectures are then discussed in Section 4. Techniques for detecting pure discrete spectrum are reviewed in Section 5 and partial results related to the Pisot Substitution Conjectures are listed in Section 6. In the final Section 7 we propose extensions of the conjectures to higher-dimensional substitutions.

2. Substitutions

In this section we review basic material concerning the dynamical systems generated by substitutions. For more detail on symbolic systems see [96, 95]; standard references for tiling systems are [113, 13, 99, 100].

2.1. Symbolic substitutions

Let $\mathcal{A}$ be a finite set, called the *alphabet*, usually $\mathcal{A} = \{1, \ldots, m\}$; its elements are called *symbols* or *letters*. Endowed with the concatenation of words as product

operation, the set $\mathcal{A}^*$ of all finite words over $\mathcal{A}$ is the *free monoid* generated by $\mathcal{A}$. If a finite or infinite word u can be factored as $u = vzw$ (v or w may be empty) we say that z is a *subword* of u.

A morphism of the free monoid $\mathcal{A}^*$ is called a *substitution* on $\mathcal{A}$. Such a substitution ϕ is said to be *primitive* if there exists a positive integer k such that, for all $i, j \in \mathcal{A}$, the word $\phi^k(i)$ contains at least one occurrence of the letter j. For any word $w \in \mathcal{A}^*$, we denote by $|w|_j$ the number of occurrences of the letter j in w and by $|w| := \sum_{j\in\mathcal{A}} |w|_j$ the *length* of w. The map $f : \mathcal{A}^* \to \mathbb{N}^d,\ w \mapsto (|w|_1, |w|_2, \ldots, |w|_m)$ is called the *abelianization map*, or *Parikh mapping*. The *substitution* (or *incidence* or *abelianization*) *matrix*, denoted by M_ϕ, is the matrix whose jth column is $f(\phi(j))$: its ijth entry is $|\phi(j)|_i$, the number of occurrences of the letter i in $\phi(j)$. Note that the substitution ϕ is primitive if and only if M_ϕ is primitive (that is, some power of M_ϕ is strictly positive). If the characteristic polynomial of M_ϕ is irreducible over $\mathbb{Q}$ then ϕ is said to be *irreducible*, and if $\det M_\phi = \pm 1$, then ϕ is *unimodular*. If ϕ is primitive, the Perron–Frobenius Theorem asserts that M_ϕ has a simple positive eigenvalue λ, which we call the *PF-eigenvalue* of ϕ, that is larger than the absolute value of all other eigenvalues. The *language* of ϕ is the subset $\mathcal{L} \subset \mathcal{A}^*$ consisting of all subwords of words of the form $\phi^k(i)$, $i \in \mathcal{A}$, $k \in \mathbb{N}$; the elements of $\mathcal{L}$ are called *admissible words* for ϕ.

To associate a symbolic dynamical system with a substitution ϕ, we give $\mathcal{A}$ the discrete topology, endow $\mathcal{A}^{\mathbb{Z}}$ with the corresponding product topology and 'extend' ϕ to $\mathcal{A}^{\mathbb{Z}}$ by $\phi((\cdots u_{-1} \cdot u_0 u_1 \cdots)) = (\cdots \phi(u_{-1}) \cdot \phi(u_0)\phi(u_1) \cdots)$, where $\cdot$ indicates the location of the 0th letter. A *substitution-periodic point* for ϕ, or *ϕ-periodic point*, is a point $u = (u_n)_{n\in\mathbb{Z}} \in \mathcal{A}^{\mathbb{Z}}$ for which $\phi^k(u) = u$ for some $k > 0$, and has the property that $u_{-1}u_0$ belongs to $\mathcal{L}$. Let s stand for the shift on $\mathcal{A}^{\mathbb{Z}}$, *i.e.*, $s((u_n)_{n\in\mathbb{Z}}) = (u_{n+1})_{n\in\mathbb{Z}}$. For a ϕ-periodic point u, let $\overline{\mathcal{O}(u)}$ be the orbit closure of u under the action of the shift s, *i.e.*, the closure in $\mathcal{A}^{\mathbb{Z}}$ of the set $\mathcal{O}(u) = \{s^j(u) : j \in \mathbb{Z}\}$.

It is not hard to see that every primitive substitution has at least one substitution-periodic point. Furthermore, the collection of all finite subwords of any substitution-periodic point of a primitive substitution equals $\mathcal{L}$ ([96]). Hence, if ϕ is a primitive substitution, then $X_\phi := \overline{\mathcal{O}(u)}$ does not depend on the choice of the ϕ-periodic point u: X_ϕ is the space of all bi-infinite words all of whose finite subwords belong to $\mathcal{L}$. The elements of X_ϕ are called *allowed words* for ϕ. The set X_ϕ is a closed, shift-invariant subset of $\mathcal{A}^{\mathbb{Z}}$ and the subshift obtained by restricting the shift s to X_ϕ is denoted by (X_ϕ, s) and called the *symbolic dynamical system*, or *substitutive system, associated with* ϕ. The system (X_ϕ, s) is *minimal* (every non-empty closed shift-invariant subset equals the whole set) and *uniquely ergodic* (there is a unique, ergodic, shift-invariant Borel probability measure on X_ϕ – see [96] for more detail).

2.2. Tile substitutions

We now consider substitutions acting on tiles, rather than letters. General references on this subject are [81, 99, 100, 114, 120].

We begin with a set of types (or labels, or colors) $\{1, \ldots, m\}$. A *tile* in $\mathbb{R}^n$ is a pair $T = (A, i)$ where $A = \text{supp}(T)$, the *support of* T, is a compact set in $\mathbb{R}^n$ which is the closure of its interior, and $i = l(T) \in \{1, \ldots, m\}$ is the *type of* T. For $x \in \mathbb{R}^n$, define $T - x := (A - x, i)$. We say that a set $\mathcal{P}$ of tiles is a *patch* if the number of tiles in $\mathcal{P}$ is finite and the tiles of $\mathcal{P}$ have mutually disjoint interiors. A tiling of $\mathbb{R}^n$ is a set $\mathcal{T}$ of tiles such that $\mathbb{R}^n = \cup_{T\in\mathcal{T}}\text{supp}(T)$ and distinct tiles have disjoint interiors. We assume that any two $\mathcal{T}$-tiles of the same type are translationally equivalent, hence there are finitely many $\mathcal{T}$-tiles up to translation.

A tiling $\mathcal{T}$ is said to have *finite local complexity* (*FLC*) if for each $R > 0$ there are, up to translation, only finitely many distinct patches in $\mathcal{T}$ with support of diameter less than $2R$ and $\mathcal{T}$ is said to be *repetitive* if for each patch $\mathcal{P} \subset \mathcal{T}$, translates of $\mathcal{P}$ occur with bounded gap in $\mathcal{T}$, *i.e.*, the set $\{x : \mathcal{P} - x \subset \mathcal{T}\}$ is relatively dense in $\mathbb{R}^n$.

Let Λ be an expanding linear map of $\mathbb{R}^n$ (meaning all eigenvalues of Λ have absolute value greater than 1) and let $\mathcal{A} = \{T_1, \ldots, T_m\}$, $T_i = (A_i, i)$, be a finite collection of tiles in $\mathbb{R}^n$. The T_i will be called *prototiles*. Let $\mathcal{A}^+$ denote the collection of all patches made of translates of prototiles. We say that $\Phi : \mathcal{A} \to \mathcal{A}^+$ is a *tile substitution* (or simply a *substitution*) with *expansion map* Λ if there are finite sets $\mathcal{D}_{ij} \subset \mathbb{R}^n$, $1 \le i, j \le m$, called *digit sets*, such that

$$\Phi(T_j) = \{T_i + v : \ v \in \mathcal{D}_{ij}, \ i = 1, \ldots, m\}, \quad \text{for } 1 \le j \le m, \tag{2.1}$$

with

$$\Lambda A_j = \bigcup_{i=1}^{m}(A_i + \mathcal{D}_{ij}) = \bigcup_{i=1}^{m}\bigcup_{d\in\mathcal{D}_{ij}}(A_i + d), \tag{2.2}$$

where the sets in the last unions of (2.2) have disjoint interiors (some of the $\mathcal{D}_{ij}$ may be empty).

When distinct tiles have translationally inequivalent supports, labeling is not necessary to distinguish tiles and we may simply identify a tile with its support. In this case, Equation (2.2) alone suffices to define a tile substitution.

The substitution (2.1) is extended to translates of prototiles by $\Phi(T_j - x) = \Phi(T_j) - \Lambda x$, and to patches and tilings by $\Phi(\mathcal{P}) = \cup_{T\in\mathcal{P}}\Phi(T)$ and $\Phi(\mathcal{T}) = \cup_{T\in\mathcal{T}}\Phi(T)$. If $\Phi(\mathcal{T}) = \mathcal{T}$, $\mathcal{T}$ is said to be *a fixed point* of the substitution Φ and $\mathcal{T}$ is Φ-*periodic* if $\Phi^k(\mathcal{T}) = \mathcal{T}$ for some $k \ge 0$. The substitution Φ may be iterated, producing larger and larger patches $\Phi^k(T_j - x)$, and by taking appropriate limits, tilings of $\mathbb{R}^n$.

As for symbolic substitutions, we associate with Φ its $m \times m$ *substitution matrix* M_Φ with ijth entry $\sharp\mathcal{D}_{ij}$. The substitution Φ is said to be *primitive* if M_Φ is primitive. A patch $\mathcal{P}$ is *admissible for* Φ if there are k, j and x so that $\mathcal{P} \subset \Phi^k(T_j - x)$. The *substitution tiling space associated with* Φ is the collection Ω_Φ of all tilings of $\mathbb{R}^n$ each of whose patches is admissible for Φ: such tilings are said to be *allowed for* Φ. There is a natural metric topology (generated by the *tiling metric*) on Ω_Φ in which two tilings are close if one agrees exactly with a small translate of the other in a large neighborhood of the origin in $\mathbb{R}^n$. Clearly,

$\mathbb{R}^n$ acts on Ω_Φ by translation and if Φ is primitive, and any $\mathcal{T} \in \Omega_\Phi$ has FLC, then the dynamical system $(\Omega_\Phi, \mathbb{R}^n)$ is compact, connected, minimal, and uniquely ergodic. In this latter case, let μ be the unique invariant Borel probability measure for the $\mathbb{R}^n$-action; then we get a measure-preserving system $(\Omega_\mathcal{T}, \mathbb{R}^n, \mu)$. Moreover, when $\mathcal{T}$ is non-periodic (meaning $\mathcal{T} - x = \mathcal{T} \implies x = 0$), $\Phi : \Omega_\Phi \to \Omega_\Phi$ is a homeomorphism that interacts with the $\mathbb{R}^n$-action via $\Phi(\mathcal{T} - x) = \mathcal{T} - \Lambda x$ ([114]). For more on substitution tilings and associated dynamical systems, see [99, 113].

We will occasionally consider a space associated with a single tiling $\mathcal{T}$ of $\mathbb{R}^n$. The *hull* of $\mathcal{T}$ is the set $\Omega_\mathcal{T} := \overline{\{\mathcal{T} - v : v \in \mathbb{R}^n\}}$, the closure being taken with respect to the tiling metric. If Φ is a primitive tile substitution and $\mathcal{T} \in \Omega_\Phi$ has FLC, then $\Omega_\mathcal{T} = \Omega_\Phi$ by minimality of the $\mathbb{R}^n$-action. When we say that an FLC tiling $\mathcal{T}$ is a *primitive substitution tiling*, we will mean that $\mathcal{T} \in \Omega_\Phi$ for some primitive tile substitution Φ. In particular, such $\mathcal{T}$ are repetitive.

From symbolic substitutions to tile substitutions. A symbolic substitution naturally gives rise to a one-dimensional tile substitution as follows. Given a primitive substitution ϕ on the alphabet $\mathcal{A} = \{1, \dots, m\}$, let $\ell = (\ell_1, \dots, \ell_m)$ be a positive left eigenvector of M_ϕ for the PF-eigenvalue λ. The prototiles for the associated tile substitution Φ are the labeled intervals $T_i := ([0, \ell_i], i)$, $i = 1, \dots, m$, and Φ is given by $\Phi(T_j) := \{T_{j_k} + \sum_{i=1}^{k-1} \ell_{j_i} : k = 1, \dots, |\phi(j)|\}$, where j_i denotes the ith letter of $\phi(j)$. Effectively, the jth prototile is stretched by a factor of λ and subdivided into translates of prototiles following the pattern $\phi(j)$. (The set $\mathcal{D}_{ij}$ of (2.1) consists of all $v = \sum_{s=1}^{k-1} \ell_{j_s}$ for which $j_k = i$, so the matrix M_Φ for the tile substitution equals the matrix M_ϕ for the symbolic substitution ϕ.) The elements of the substitution tiling space $\Omega_\phi := \Omega_\Phi$ are tilings of the real line whose tiles follow the pattern of some allowed bi-infinite word in X_ϕ. Tilings $\mathcal{T} \in \Omega_\phi$ can be thought of as interpolates of elements of X_ϕ: formally, the system $(\Omega_\phi, \mathbb{R})$ is topologically conjugate with the suspension system $(\hat{X}_\phi, \mathbb{R})$ of (X_ϕ, s) with roof function given by ℓ (see, *e.g.*, [51, Chap.2] for a discussion of suspension). As ϕ is primitive, so is Φ, and $(\Omega_\phi, \mathbb{R})$ is compact, connected, minimal and uniquely ergodic. Moreover, Φ is a homeomorphism of Ω_ϕ and $\Phi(\mathcal{T} - t) = \mathcal{T} - \lambda t$ for $\mathcal{T} \in \Omega_\phi$ and $t \in \mathbb{R}$.

2.3. Point set substitutions

Point sets provide the most flexible context for studying substitutive dynamics. A *Delone set* is a relatively dense and uniformly discrete subset of $\mathbb{R}^n$. We say that $\underline{\Gamma} = \bigcup_{i=1}^m \Gamma_i \times \{i\}$ is a *Delone multi-color set* in $\mathbb{R}^n$ if each Γ_i is Delone and $\mathrm{supp}(\underline{\Gamma}) := \bigcup_{i=1}^m \Gamma_i \subset \mathbb{R}^n$ is Delone. We call such $\underline{\Gamma}$ a *substitution Delone multi-color set* if $\underline{\Gamma}$ is a Delone multi-color set and there exist an expanding map $\Lambda : \mathbb{R}^n \to \mathbb{R}^n$ and finite sets $\mathcal{D}_{ij}$ (the *digit sets*), for $1 \le i, j \le m$, such that

$$\Gamma_i = \bigcup_{j=1}^m (\Lambda \Gamma_j + \mathcal{D}_{ij}), \quad 1 \le i \le m, \tag{2.3}$$

where the union on the right-hand side is disjoint. The substitution Delone multi-color set $\underline{\Gamma}$ is said to be *primitive* if the matrix $(\sharp\mathcal{D}_{ij})_{m\times m}$ is primitive. For more on substitution Delone sets, see, *e.g.*, [74, 81].

For any given substitution Delone multi-color set $\underline{\Gamma} = \bigcup_{i=1}^{m} \Gamma_i \times \{i\}$, we define Φ_{ij} to be the collection of affine functions from $\mathbb{R}^n \times \{j\}$ to $\mathbb{R}^n \times \{i\}$

$$\Phi_{ij} := \{f : (x, j) \mapsto (\Lambda x + a, i) \,:\, a \in \mathcal{D}_{ij}\}. \tag{2.4}$$

Then $\Phi_{ij}(\Gamma_j \times \{j\}) := \bigcup_{f\in\Phi_{ij}} f(\Gamma_j \times \{j\}) = (\Lambda\Gamma_j + \mathcal{D}_{ij}) \times \{i\}$. We denote by Φ the $m \times m$ array with ijth entry Φ_{ij} and call Φ a *matrix function system* (MFS). For any $\underline{S} = \bigcup_{i=1}^{m} S_i \times \{i\}$, $S_i \subset \mathbb{R}^n$, we define $\Phi(\underline{S})$ to be the collection $\bigcup_{i=1}^{m}(\bigcup_{j=1}^{m} \Phi_{ij}(S_j \times \{j\})) \times \{i\}$. We may then iterate Φ, obtaining, for any $k \in \mathbb{N}$, $\Phi^k(\underline{\Gamma}) = \underline{\Gamma}$ and $\Phi^k(\Gamma_j \times \{j\}) = \bigcup_{1\le i\le m}(\Lambda^k\Gamma_j + (\mathcal{D}^k)_{ij}) \times \{i\}$ where

$$(\mathcal{D}^k)_{ij} = \bigcup_{1\le n_1,n_2,\ldots,n_{(k-1)}\le m} (\mathcal{D}_{in_1} + \Lambda\mathcal{D}_{n_1n_2} + \cdots + \Lambda^{k-1}\mathcal{D}_{n_{(k-1)}j}).$$

A *cluster* of $\underline{\Gamma}$ is a collection $\underline{P} = \bigcup_{i=1,\ldots,m} P_i \times \{i\}$ where $P_i \subset \Gamma_i$ is finite for all $1 \le i \le m$. We say that a cluster $\underline{P}$ is *legal* if it is a translate of a subcluster of a cluster generated from one point of $\underline{\Gamma}$, *i.e.*, $a + \underline{P} \subset \Phi^k((x, i))$ for some $k \in \mathbb{N}$, $a \in \mathbb{R}^n$ and $x \in \Gamma_i$.

As for tilings, we may consider the hull $\Omega_{\underline{\Gamma}}$ of a Delone multi-color set and the associated dynamical system $(\Omega_{\underline{\Gamma}}, \mathbb{R}^n)$.

The equations (2.3) and (2.2) are formally the same. Indeed, one may pass back and forth between substitution tilings and substitution Delone multi-color sets with a fair amount of freedom as we now describe.

From a substitution tiling to a substitution Delone multi-color set. From a given substitution tiling which is a fixed point of a substitution, one easily constructs a substitution Delone multi-color set by taking a representative point from each tile, choosing points in the same relative position in tiles of the same color. In fact, if $\mathcal{T}$ is a repetitive tiling that is a fixed point of a substitution for which $\mathcal{T} = \bigcup_{i=1}^{m}\{T_i + x : x \in \Gamma_i\}$, then $\underline{\Gamma} = \bigcup_{i=1}^{m} \Gamma_i \times \{i\}$ is a primitive substitution Delone multi-color set. We call $\underline{\Gamma}_{\mathcal{T}}$ an associated substitution Delone multi-color set of $\mathcal{T}$.

A particularly nice selection of associated Delone multi-color set is as follows. For each $T \in \mathcal{T}$, let $\tau(T) \in \Phi(T)$ be chosen in such a way that if $S = T + x$ then $\tau(S) = \tau(T)+\Lambda x$. Since we are assuming that $\mathcal{T}$ is a fixed point of the substitution, $\tau : \mathcal{T} \to \mathcal{T}$. Let

$$\Gamma_i := \bigcup_{T\in\mathcal{T},\, T \text{ of type } i} \bigcap_{k=0}^{\infty} \Lambda^{-k}\mathrm{supp}(\tau^k(T)).$$

For Γ_i so defined, $\mathcal{C} := \bigcup_{i=1}^{m} \Gamma_i$ is called a set of *control points* for $\mathcal{T}$.

From a substitution Delone multi-color set to a substitution tiling. On the other hand, it is not so obvious how to reconstruct a substitution tiling from a substitution Delone multi-color set. Lagarias and Wang give a canonical way to carry out such a construction in a restricted setting (see [74, 81] for more detail). We briefly describe the context in which this construction applies. We say that a Delone multi-color set $\underline{\Gamma} = \bigcup_{i=1}^{m} \Gamma_i \times \{i\}$ is *representable* (by tiles) if there exist tiles $T_i = (A_i, i), 1 \le i \le m$, so that

$$\{x + T_i : \ x \in \Gamma_i, 1 \le \ i \le m\} \quad \text{is a tiling of } \ \mathbb{R}^n, \tag{2.5}$$

that is, $\mathbb{R}^n = \bigcup_{1 \le i \le m} \bigcup_{x \in \Gamma_i} (x + A_i)$, and the sets in this union have disjoint interiors. In the case that $\underline{\Gamma}$ is a primitive substitution Delone multi-color set, we will understand the term representable to mean relative to tiles $T_i = (A_i, i), 1 \le i \le m$, that satisfy the adjoint equations (2.2) formed by the digit sets $\mathcal{D}_{ij}$. In [74, Lemma 3.2] it is shown that if $\underline{\Gamma}$ is a substitution Delone multi-color set, then there is a finite multi-color set (cluster) $\underline{P} \subset \underline{\Gamma}$ for which $\Phi^{n-1}(\underline{P}) \subset \Phi^n(\underline{P})$ for $n \ge 1$ and $\underline{\Gamma} = \lim_{n \to \infty} \Phi^n(\underline{P})$. We call such a multi-color set $\underline{P}$ a *generating set* for $\underline{\Gamma}$. It is shown in [81] that if a generating multi-color set of a primitive substitution Delone multi-color set $\underline{\Gamma}$ is legal, then $\underline{\Gamma}$ is representable.

2.4. Beta-shifts

The final context for substitution dynamics that we consider arises from transformations of the interval associated with beta-numeration. For more detail, the reader is referred to [85].

Let $\beta > 1$ be a real number. The *beta transformation*, $T_\beta : [0,1) \to [0,1)$, is defined by $T_\beta(x) = \beta x - d(x)$ with $d(x) = \lfloor \beta x \rfloor \in [0, \beta) \cap \mathbb{Z}$. Putting $d_n(x) = d(T^{n-1}(x))$ for $n = 0, 1, \ldots$, we obtain the *β-expansion of x*:

$$x = \frac{d_1(x)}{\beta} + \frac{d_2(x)}{\beta^2} + \frac{d_3(x)}{\beta^3} + \cdots .$$

Define a function $\mathbf{d}_\beta : [0,1] \to \mathcal{A}^{\mathbb{N}}$ by $\mathbf{d}_\beta(x) = d_1(x)d_2(x)d_3(x)\cdots$ with $\mathcal{A} = [0, \beta) \cap \mathbb{Z}$. For $x > 1$, we can find the minimum non-negative integer m with $\beta^{-m} x \in [0,1)$. Putting $d_{-m+k}(x) = d_k(\beta^{-m} x)$, we obtain:

$$x = d_{-m}(x)\beta^m + d_{-m+1}(x)\beta^{m-1} + \cdots + d_0(x) + \frac{d_1(x)}{\beta} + \frac{d_2(x)}{\beta^2} + \frac{d_3(x)}{\beta^3} + \cdots$$

which we denote by $d_{-m}(x)d_{-m+1}(x)\cdots d_0(x)d_1(x)d_2(x)d_3(x)\cdots$. It is sometimes convenient to introduce a special symbol $\bullet$, the 'decimal point', to indicate the initial position: $d_{-m}(x)d_{-m+1}(x)\cdots d_0(x) \bullet d_1(x)d_2(x)d_3(x)\cdots$ Here $\bullet$ appears only once in the expansion and we ignore $\bullet$ when we treat this bi-infinite word as an element of $\mathcal{A}^{\mathbb{Z}}$. Note that the mark here appears to the right of the 0th element.

Not every word in $\mathcal{A}^{\mathbb{N}}$ is realized as a beta expansion, as, for example, 9^∞ is not allowed in the tail of a decimal expansion. A word in $\mathcal{A}^{\mathbb{N}}$ is called *admissible* if it is a beta expansion of some $x \ge 0$. We also say a finite word $\omega \in \mathcal{A}^*$ is admissible if $\omega 0^\infty = \omega 00\ldots$ is admissible.

The *beta shift* X_β is the subshift of the full shift on $\mathcal{A}^{\mathbb{Z}}$ consisting of the collection of bi-infinite words over $\mathcal{A}$, all of whose finite subwords are admissible. The function T_β is a discontinuous piecewise linear transform on $[0,1)$ and its discontinuities are essential in describing admissibility. A nice feature of beta-expansion is that admissibility, as we now describe, is easily computed through consideration of the *expansion of one*:

$$\mathbf{d}_\beta(1-) = \lim_{\varepsilon \downarrow 0} \mathbf{d}_\beta(1-\varepsilon)$$

where the limit is in the topology of $\mathcal{A}^{\mathbb{N}}$. The expansion $\mathbf{d}_\beta(1-) = c_1 c_2 \cdots$ is made concrete in the following way. We have $c_1 = \lfloor \beta \rfloor$ and if $T^n(\beta - \lfloor \beta \rfloor) \neq 0$ for all $n \geq 0$, then $c_{n+2} = \lfloor \beta T^n(\beta - \lfloor \beta \rfloor) \rfloor$. If there is n with $T^n(\beta - \lfloor \beta \rfloor) = 0$, then take the smallest n with this property, in which case we have $\mathbf{d}_\beta(1-) = (c_1 c_2 \cdots c_n(c_{n+1} - 1))^\infty$. The admissibility condition, called the *Parry condition* is then: a word $x \in \mathcal{A}^{\mathbb{N}}$ is admissible if and only if $s^n(x) \ll \mathbf{d}_\beta(1-)$, where $\ll$ is the lexicographic order (see [98, 66]) and s denotes the shift on $\mathcal{A}^{\mathbb{N}}$. By this characterization, we see that X_β is the set of all bi-infinite sequences of labels of bi-infinite walks on a finite labeled directed graph. More precisely, X_β is a subshift of finite type if and only if $\mathbf{d}_\beta(1-)$ is purely periodic, and it is sofic if and only if $\mathbf{d}_\beta(1-)$ is eventually periodic. If X_β is sofic, then β is a Perron number, and if β is a Pisot number, then X_β must be sofic. Assume that X_β is sofic and put $\mathbf{d}_\beta(1-) = c_1 c_2 \cdots c_m (c_{m+1} c_{m+2} \cdots c_{m+\ell})^\infty$. Note that we always have $\ell > 0$, and if X_β is a subshift of finite type, then $m = 0$. Figure 1 shows examples of such graphs for X_β with $\mathbf{d}_\beta(1-)$ as above.

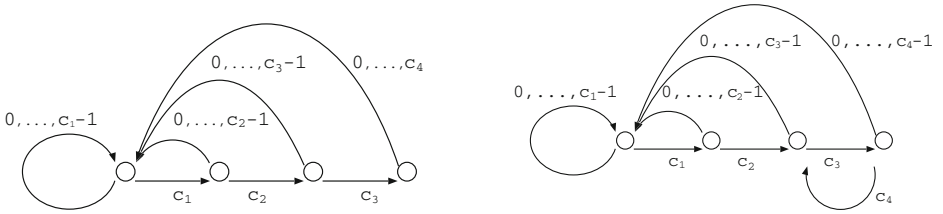

(A) $(c_1c_2c_3c_4)^\infty$ (B) $c_1c_2(c_3c_4)^\infty$

FIGURE 1. Graph of X_β

We write an element $x = (a_n)_{n \in \mathbb{Z}} \in X_\beta$ as $x = \cdots a_{-2}a_{-1}a_0 \bullet a_1 a_2 \cdots$ and say that $x_I := \cdots a_{-2}a_{-1}a_0$ (resp. $x_F := a_1 a_2 \cdots$) is the *integer part* (resp. the *fractional part*) of x. The *future* of the integer part $x_I = \cdots a_{-2}a_{-1}a_0$ is defined by:

$$\mathcal{F}(x_I) = \{b_1 b_2 \cdots \in \mathcal{A}^{\mathbb{N}} \mid \cdots a_{-2}a_{-1}a_0 \bullet b_1 b_2 \cdots \in X_\beta\}$$

and the *past* of the fractional part $x_F = a_1 a_2 \cdots$ is defined by:

$$\mathcal{P}(x_F) = \{\cdots b_{-2}b_{-1}b_0 \in \mathcal{A}^{\mathbb{N}} | \ \cdots b_{-2}b_{-1}b_0 \bullet a_1 a_2 \cdots \in X_\beta\}.$$

By the above graph characterization, it is clear that X_β is sofic if and only if there are only finitely many distinct future sets $\mathcal{F}(x_I)$ and past sets $\mathcal{P}(x_F)$. If

X_β is sofic, we may associate with β a tile substitution with expansion β and (unlabeled) prototiles $\mathcal{F}(x_I)$ by means of the natural map

$$\pi : a_1 a_2 \cdots \mapsto \sum_{i=1}^{\infty} a_i \beta^{-i} \tag{2.6}$$

that satisfies

$$\beta\pi(\mathcal{F}(x_I)) = \bigcup_{x_I a \text{ is admissible}} \pi(\mathcal{F}(x_I a)) + \pi(a\bullet) \tag{2.7}$$

where $x_I a$ denotes the concatenation of x_I and $a \in \mathcal{A}$ and $\pi(a\bullet)$ is just the integer a.

From a beta-shift to a symbolic substitution. A symbolic substitution, denoted ϕ_β, and called the *β-substitution*, may be associated with a sofic beta-shift. This β-substitution is defined by:

$$\begin{aligned} \phi_\beta(1) &= 1^{c_1}2 \\ \phi_\beta(2) &= 1^{c_2}3 \\ &\vdots \\ \phi_\beta(m) &= 1^{c_m} \end{aligned}$$

for the shift of finite type case with $\mathbf{d}_\beta(1-) = (c_1 \cdots c_m)^\infty$,

$$\begin{aligned} \phi_\beta(1) &= 1^{c_1}2 \\ \phi_\beta(2) &= 1^{c_2}3 \\ &\vdots \\ \phi_\beta(m+\ell-1) &= 1^{c_{m+\ell-1}}(m+\ell) \\ \phi_\beta(m+\ell) &= 1^{c_{m+\ell}}(m+1) \end{aligned}$$

for the general sofic case.

It is easy to see that ϕ_β is primitive, although it may not be irreducible. By primitivity, we get minimal and uniquely ergodic substitutive symbolic and tiling dynamical systems, (X_{ϕ_β}, s) and $(\Omega_{\phi_\beta}, \mathbb{R})$. Up to a rescaling of the $\mathbb{R}$-action and the labeling of tiles, the system $(\Omega_{\phi_\beta}, \mathbb{R})$ is just the system derived from the tiling substitution (2.7).

3. Discreteness of the dynamical and diffraction spectra

3.1. Eigenvalues and coincidence rank

Recall that an eigenfunction for an $\mathbb{R}^n$-action on a space Ω with invariant measure μ is an $L^2(\Omega, \mu)$ function[1] f for which there is an associated eigenvalue $\alpha \in \mathbb{R}^n$ so that $f(T - v) = e^{2\pi i \langle \alpha, v \rangle} f(T)$ for all $v \in \mathbb{R}^n$ and $T \in \Omega$. The $\mathbb{R}^n$-action is said to have *pure discrete spectrum*[2] or *pure point spectrum* if the linear span of the eigenfunctions is dense in $L^2(\Omega, \mu)$. When the $\mathbb{R}^n$-action is ergodic, eigenfunctions must have constant absolute value and eigenvalues are simple: if f and g are eigenfunctions with eigenvalue α, then $f = cg$ for some $c \in \mathbb{C}$. In this situation we may choose eigenfunctions to have values in the unit circle $\mathbb{T}^1$. Furthermore, if $\Omega = \Omega_\Phi$

[1] In this measure-theoretical framework, statements about eigenfunctions should be interpreted to hold a.e.

[2] Note that here 'discrete' refers to a property of the span of the eigenfunctions and not to discreteness of the eigenvalues as a set.

is a substitution tiling space, eigenfunctions can be chosen to be continuous ([115]). We may view each continuous eigenfunction as a semi-conjugacy, or factor map, between $(\Omega, \mathbb{R}^n)$ and an action of $\mathbb{R}^n$ by translation on the compact abelian group $\mathbb{T}^1$.

The definitions for a $\mathbb{Z}$-action on a space are analogous: for example, an eigenfunction f for the shift s on a symbolic space X with associated eigenvalue $\alpha \in \mathbb{R}$ satisfies $f(s^n(x)) = e^{2\pi i \alpha n} f(x)$ for all $n \in \mathbb{Z}$ and $x \in X$.

It is a consequence of the Halmos-von Neumann theory that the $\mathbb{R}^n$-action on a tiling space (likewise, the shift on a symbolic space) has pure discrete spectrum if and only if the action is measurably conjugate to translation on a compact abelian group (see [121, 96], and see Theorems 5.2 and 7.8 for illustrations).

Every primitive FLC substitution tiling dynamical system $(\Omega_\Phi, \mathbb{R}^n)$ has a maximal equicontinuous factor $(\hat{\mathbb{T}}, \mathbb{R}^n)$ with factor map $g : \Omega \to \hat{\mathbb{T}}$ and g is a.e. m-to-1 for some $m \in \mathbb{N} \cup \{\infty\}$. (Here $\hat{\mathbb{T}}$ is a torus or solenoid, the $\mathbb{R}^n$-action is a Kronecker action, and the map g is obtained by considering all eigenfunctions.) The number m is called the *coincidence rank* of Φ and is denoted by $cr(\Phi)$: by the Halmos-von Neumann theory, $(\Omega_\Phi, \mathbb{R}^n)$ has pure discrete spectrum if and only if $cr(\Phi) = 1$. It follows from [113, 29, 21] that, for one-dimensional primitive tile substitutions Φ, $cr(\Phi) < \infty$ if and only if the expansion λ for Φ is a Pisot number.

If the dynamical spectrum of a tiling $\mathcal{T}$ is pure discrete, the eigenvalues for the dynamical system $(\Omega_\mathcal{T}, \mathbb{R}^n)$ span $\mathbb{R}^n$. Since every additive combination of eigenvalues for the dynamical system is also an eigenvalue, the eigenvalues are relatively dense. In the next section we translate this necessary condition for pure discrete spectrum (that is, relative denseness of the eigenvalues) into a condition expressed in terms of substitution Delone multi-color sets, namely, the Meyer property.

3.2. The Meyer property

For the study of the discreteness of the dynamical spectrum of a tiling space, the so-called Meyer property plays an important role. A Delone set Y is called a *Meyer set* if $Y - Y$ is uniformly discrete. Necessarily, if Y is a Meyer set, then Y has FLC. Suppose that $\mathcal{T}$ is a primitive FLC substitution tiling in $\mathbb{R}^n$ and that $\underline{\varGamma}_\mathcal{T} = \bigcup_{i=1}^m \varGamma_i \times \{i\}$ is a substitution Delone multi-color set associated with $\mathcal{T}$, as described in Section 2.3. When $\varGamma_\mathcal{T} = \cup_{1 \le i \le m} \varGamma_i$ is a Meyer set, we say that the tiling $\mathcal{T}$ has the *Meyer property*. For characterizations of the Meyer property, see [73] for $\mathbb{R}^n$, and [20] for the general case of a σ-compact, locally compact abelian group.

The connection between Meyer sets and the spectrum of an FLC substitution tiling can be viewed through the mathematical concept of diffraction measure developed by Hof [61, 62]. For a Delone set $\varGamma$, consider $\delta_\varGamma = \sum_{x \in \varGamma} \delta_x$. Let $\gamma(\delta_\varGamma)$ denote its autocorrelation (assuming it is unique), that is, the vague limit[3]

$$\gamma(\delta_\varGamma) = \lim_{r \to \infty} \frac{1}{\mathrm{Vol}(B_r)} \left(\delta_\varGamma|_{B_r} * \widetilde{\delta_\varGamma|_{B_r}} \right), \tag{3.1}$$

[3]Recall that if f is a function in $\mathbb{R}^n$, then $\tilde{f}$ is defined by $\tilde{f}(x) = \overline{f(-x)}$. If μ is a measure, $\tilde{\mu}$ is defined by $\tilde{\mu}(f) = \overline{\mu(\tilde{f})}$ for all $f \in \mathcal{C}_0(\mathbb{R}^n)$.

where B_r is a ball of radius r. The Fourier transform $\widehat{\gamma(\delta_\Gamma)}$ is called the *diffraction measure* for δ_Γ. We say that the measure δ_Γ, or the Delone set Γ has *pure point diffraction spectrum*, if $\widehat{\gamma(\delta_\Gamma)}$ is a pure point or discrete measure. The point masses of the diffraction measure are called *Bragg peaks*.

There is a close correspondence between pure pointedness of the diffraction measure and pure discrete spectrum of the associated dynamical system. For example, if $\mathcal{T}$ is a primitive FLC substitution tiling in $\mathbb{R}^n$ and $\underline{\Gamma}_{\mathcal{T}} = \bigcup_{i=1}^m \Gamma_i \times \{i\}$ is a substitution Delone multi-color set associated with $\mathcal{T}$, then each Γ_i has pure point diffraction spectrum if and only if $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ has pure discrete spectrum ([17, 18, 48, 60, 80, 81, 82, 102]).

The following theorem is due to Strungaru [116] and Lee–Solomyak [77].

Theorem 3.1 ([116, 77]). *Suppose that $\mathcal{T}$ is a primitive FLC substitution tiling in $\mathbb{R}^n$ and that $\underline{\Gamma}_{\mathcal{T}} = \bigcup_{i=1}^m \Gamma_i \times \{i\}$ is a substitution Delone multi-color set associated with $\mathcal{T}$. The following are equivalent:*

(i) *The set of locations of Bragg peaks for each Γ_i is relatively dense in $\mathbb{R}^n$.*
(ii) *The set of eigenvalues for $(\Omega_{\mathcal{T}}, \mathbb{R}^n, \mu)$ is relatively dense in $\mathbb{R}^n$.*
(iii) *$\Gamma = \cup_{1 \le i \le m} \Gamma_i$ is a Meyer set.*

In particular, if $\mathcal{T}$ is a primitive FLC substitution tiling for which the dynamical system $(\Omega_{\mathcal{T}}, \mathbb{R}^n, \mu)$ has pure discrete spectrum, then $\Gamma_{\mathcal{T}} = \cup_{1 \le i \le m} \Gamma_i$ is a Meyer set. In other words, the Meyer property is necessary for pure discrete spectrum of an FLC substitution tiling space. This is not generally true for the hull of a tiling that does not arise from a substitution: it fails, for example, for modulated lattices (see [16, 79, 110, 118]).

3.3. Cut-and-project schemes and model sets

We now endeavor to fill the gap between relative density for eigenvalues and pure discreteness of the dynamical spectrum. This is the aim of the present section where the notions of regular model set and inter-model set are introduced in the setting of cut-and-project schemes. For more about model sets see, *e.g.*, the surveys [88, 89] and [16, Chap. 7] where an explicit discussion with the silver mean as an example can be found.

A *cut-and-project scheme* (or *CPS* for short) consists of a triple $(G, H, \widetilde{L})$ where G and H are σ-compact, locally compact abelian groups, $\widetilde{L}$ is a lattice, *i.e.*, a discrete subgroup for which the quotient group $(G \times H)/\widetilde{L}$ is compact, such that the restriction of the canonical projection $\pi_1 : G \times H \to G$ to $\widetilde{L}$ is injective and the image $L^\star = \pi_2(\widetilde{L})$ of $\widetilde{L}$ under the canonical projection $\pi_2 : G \times H \to H$ is dense in H. Schematically, we summarize the situation as follows by setting $L = \pi_1(\widetilde{L})$:

$$\begin{array}{ccccc} G & \xleftarrow{\pi_1} & G \times H & \xrightarrow{\pi_2} & H \\ \cup & & \cup & & \cup \text{ dense} \\ L & \xleftarrow{1-1} & \widetilde{L} & \longrightarrow & L^\star \end{array} \tag{3.2}$$

For a subset $W \subset H$, called the *acceptance window*, let

$$\varGamma(W) = \{\pi_1(x) \in G \,:\, x \in \widetilde{L}, \pi_2(x) \in W\}\,.$$

A subset of G of the form $\varGamma = \varGamma(W) + g$, where $g \in G$ and $W \subset H$ has non-empty interior and compact closure, is called a *model set*. The non-empty interior property yields that any model set is relatively dense in G ([109]). Intuitively, being a model set means that all the points in $\varGamma$ are obtained through the selection process by the window W.

Models sets are deeply connected to Meyer sets. In fact all model sets are Meyer sets ([78]). The converse reads as follows: a relatively dense set in G is a Meyer set if and only if it is a subset of a model set of G ([88]). See [19, 20, 76, 89, 101] for more on model sets.

Applying Theorem 3.1, we obtain the following.

Corollary 3.2. *Suppose that $\mathcal{T}$ is a primitive substitution tiling in $\mathbb{R}^n$ and that $\underline{\varGamma}_{\mathcal{T}} = \bigcup_{i=1}^{m} \varGamma_i \times \{i\}$ is the substitution Delone multi-color set associated with $\mathcal{T}$. If $\varGamma = \cup_{1\le i\le m}\varGamma_i$ is a model set, then the set of eigenvalues for $(\Omega_{\mathcal{T}}, \mathbb{R}^n, \mu)$ is relatively dense in $\mathbb{R}^n$.*

Note that additional conditions are needed to guarantee pure discrete spectrum. We say that a model set $\varGamma$ is *regular* if the boundary $\partial W = \overline{W} \setminus W^\circ$ of W is of (Haar) measure 0.

Theorem 3.3 ([102]). *If $\varGamma \subset R^n$ is a regular model set, then $(\Omega_{\varGamma}, \mathbb{R}^n)$ has pure discrete spectrum.*

Theorem 3.4 ([81]). *Let $\mathcal{T}$ be a primitive substitution tiling in $\mathbb{R}^n$ with associated Delone multi-color set $\underline{\varGamma}_{\mathcal{T}} = \bigcup_{i=1}^{m}(\varGamma_i \times \{i\})$ and suppose that $L = \bigcup_{1\le i\le m} \varGamma_i$ is a lattice. Then each $\varGamma_i$ is a regular model set for $1 \le i \le m$ if and only if $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ has pure discrete spectrum.*

In the general case, we define an *inter-model set* as a subset $\varGamma$ of G for which $s + \varGamma(W^\circ) \subset \varGamma \subset s + \varGamma(W)$ for some $s \in \mathbb{R}^n$, where W is compact in H and $W = \overline{W^\circ} \neq \emptyset$. Inter-model sets are Delone sets, and all regular model sets also are inter-model sets (see [109] for more detail).

Theorem 3.5 ([75]). *Let $\mathcal{T}$ be a primitive substitution tiling in $\mathbb{R}^n$ with FLC and $\underline{\varGamma}_{\mathcal{T}} = \bigcup_{i=1}^{m}(\varGamma_i \times \{i\})$ be an associated substitution Delone multi-color set. Then $\varGamma_i$ is an inter-model set if and only if $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ has pure discrete spectrum.*

Dynamical systems generated by regular or inter-model sets coming from primitive substitutions thus have pure discrete spectrum. Sections 3.4 and 3.5 exhibit natural candidates for acceptance windows and cut-and-project schemes to be associated with symbolic substitutions and beta-numeration. Their particular structure allows a convenient formulation of sufficient conditions ensuring that these constructions indeed yield regular model sets. As will be discussed in Section 5, pure discrete substitution tiling systems can be viewed as inter-model sets. This will require the introduction of intermediary conditions based on notions of

coincidences (see in particular Theorems 5.4 and 5.5). Checking whether these additional conditions are satisfied is at the core of the Pisot Substitution Conjecture.

3.4. Rauzy fractals and symbolic substitutions

In the case of a symbolic substitution ϕ, there is a natural way to exhibit a cut-and-project scheme for the associated substitution tiling Φ. This approach was initiated by Rauzy in the seminal paper [97] and is based on the notion of Rauzy fractal, an object that will be our candidate both for being an acceptance window for a cut-and-project scheme such as defined in (3.2), and equivalently, for being a fundamental domain for the Kronecker action in case of pure discrete spectrum.

Suppose that ϕ is a symbolic Pisot substitution, and, for ease of exposition, that ϕ is also unimodular and irreducible. Thus, if d is the degree of the PF-eigenvalue, λ, of M_ϕ (d is also equal to the cardinality of the alphabet), then M_ϕ is $d \times d$ and $\mathbb{R}^d$ splits as the direct sum $\mathbb{R}^d = E^s \oplus E^u$ of $d-1$ and one-dimensional M_ϕ-invariant subspaces. Let $\pi_s : \mathbb{R}^d \to E^s$ and $\pi_u : \mathbb{R}^d \to E^u$ denote the corresponding projections and let r be a positive right PF-eigenvector for M_ϕ (so $E^u = \{t\, r : t \in \mathbb{R}\}$). Let u be a ϕ-periodic point. The *ith Rauzy piece* for ϕ is the subset of E^s

$$\mathcal{R}_i := \overline{\{\pi_s(f(u_0 \cdots u_j)) : j \geq 0, u_{j+1} = i\}}$$

and the *Rauzy fractal* for ϕ is the union of the pieces

$$\mathcal{R} := \cup_{i=1}^{d} \mathcal{R}_i.$$

The $\mathcal{R}_i$ do not depend on u, each is the closure of its interior and has zero measure boundary. For general properties of Rauzy fractals, see, *e.g.*, [97, 14, 95, 35, 39].

As constructed above, the Rauzy fractal is obtained as the closure of the stable projection of a certain subset of points from the integer lattice $\mathbb{Z}^d$. We will see in Section 3.5 below, in the context of beta-numeration, that the Rauzy fractal can also be obtained using the Minkowski embedding of $\mathbb{Q}(\lambda)$ in $\mathbb{R}^d$. This is the approach of [117] and has been extended through the use of non-Euclidean representation spaces to allow for a definition of the Rauzy fractal when the Pisot number λ (the PF-eigenvalue) is not a unit (see [107, 109]).

Rauzy fractals and cut-and-project schemes. Given an irreducible unimodular Pisot substitution ϕ on d letters, the d entries of left and right PF-eigenvectors of the substitution matrix M_ϕ are linearly independent over $\mathbb{Q}$. It follows that the projection $\pi_u : \mathbb{R}^d = E^s \oplus E^u \to E^u \simeq \mathbb{R}$ is one-to-one on the lattice $\mathbb{Z}^d$ and that $\pi_s(\mathbb{Z}^d)$ is dense in $E^s \simeq \mathbb{R}^{d-1}$. Using the Rauzy fractal $\mathcal{R} \subset E^s$ as an acceptance window, we may define a Delone set $\Gamma(\mathcal{R}) \subset E^u$ by projecting into E^u, via π_u, those elements of $\mathbb{Z}^d$ that lie in $\pi_s^{-1}(\mathcal{R})$:

$$\Gamma(\mathcal{R}) := \{\pi_u(x) \in E^u : x \in \mathbb{Z}^d,\, \pi_s(x) \in \mathcal{R}\}.$$

Let $V := \{f(u_0 \cdots u_{k-1}) : k \in \mathbb{N}\} \cup \{-f(u_{-k} \cdots u_{-1}) : k \in \mathbb{N}\}$. (The set V is the set of vertices of the *strand* S_u to be defined in Section 5.2 below.) One sees easily from the definition of $\mathcal{R}$ that $\pi_u(V) \subset \Gamma(\mathcal{R})$.

Delone sets that arise from such a cut-and-project procedure have particularly nice spectral properties and it thus becomes relevant for the study of the dynamical systems (X_ϕ, s) and $(\Omega_\phi, \mathbb{R})$ to ask the question: is $\pi_u(V)$ essentially the same as $\Gamma(\mathcal{R})$? By observing that the boundary of $\mathcal{R}$ has zero measure, a positive answer yields that (X_ϕ, s) has pure discrete spectrum. Hence, most conditions ensuring pure discrete spectrum expressed in terms of Rauzy fractals can be reformulated as sufficient conditions for the associated substitution Delone multi-color set to be a regular model set with acceptance window given by the Rauzy fractal $\mathcal{R}$. As an illustration, the equality between $\pi_u(V)$ and $\Gamma(\mathcal{R})$ can be proved in particular cases by studying the way that particular polygonal approximations of the boundary of $\mathcal{R}$ converge to this boundary ([52, 65, 84]).

Periodic multi-tilings in $\mathbb{R}^d$. In order to make explicit the fact that the Rauzy fractal is a natural candidate for a fundamental domain for the Kronecker action for the maximal equicontinuous factor of Ω_ϕ, we endeavor to construct tilings of E^s by translates of the $\mathcal{R}_i$. For each $i = 1, \ldots, d$, let e_i denote the ith standard unit vector. Let us fix a vector $y \in E^u$. Let $t_i \in \mathbb{R}^+$ be such that $t_i\, y = \pi_u(e_i)$, and let $\mathcal{C}_i := \{x + t\, y : x \in \mathcal{R}_i \text{ and } 0 \le t \le t_i\}$. The collection $\mathcal{M} := \{\mathcal{C}_i + v : 0 \le i \le d \text{ and } v \in \mathbb{Z}^d\}$ is a periodic *multi-tiling* of $\mathbb{R}^d$ by translations of the $\mathcal{C}_i$: there is $m \in \mathbb{N}$, called the *degree*, or *multiplicity* of the multi-tiling so that almost every $x \in \mathbb{R}^d$ lies in exactly m tiles of $\mathcal{M}$. In fact, the degree m is equal to the coincidence rank $cr(\phi)$ of ϕ introduced in Section 3.1 (see [30]).

Theorem 3.6 ([65, 30]). *Let ϕ be a primitive Pisot substitution that is also unimodular and irreducible. The spaces (X_ϕ, s) and $(\Omega_\phi, \mathbb{R})$ have pure discrete spectrum if and only if the degree m of the multi-tiling associated with ϕ is* 1.

(For ϕ primitive, Pisot, and irreducible, pure discreteness of (X_ϕ, s) and $(\Omega_\phi, \mathbb{R})$ are equivalent ([45]) – see Section 4.)

Several approaches have been developed in order to make the tiling condition of Theorem 3.6 algorithmic. Indeed, the intersections between the sets $\mathcal{C}_i + v$ have a self-similar structure that can be described in terms of finite graphs (see the monograph [108] and the bibliography therein for more detail). Another approach consists in studying the *dual multi-tiling* associated with ϕ, that is, the multi-tiling of E^s given by $\{T \cap E^s : T \in \mathcal{M}\}$. This dual multi-tiling of the $(d-1)$-dimensional space E^s can be described via a discrete combinatorial action, namely a generalized (also called dual) substitution ([14, 30]). Each prototile of this dual multi-tiling can be characterized by a pair made of a Rauzy piece and a face of the discrete plane associated with the contracting space E^s. This discrete plane is stable under the action of the associated dual substitution. As a consequence, further combinatorial formulations of the pure discrete spectrum property can be stated in effective terms (see [14, 39] for more detail). These conditions are dual versions of the notions of geometric coincidence developed in Section 5.2 (where the tilings live in the unstable space E^u). Observe also that the terminology 'dual' is here consistent with the notion of star-duality for self-similar cut-and-project

tilings such as developed in [117, 55]; see also [37] for an illustration of these connections in the two-letter case.

3.5. Cut-and-project schemes and beta-numeration

We present here an arithmetical version of the previous construction for the beta shift due originally to Thurston [117], which was itself inspired by Rauzy's construction [97]. For more on these constructions, see [1, 3, 5, 58].

Let β be a Pisot unit of degree d with real algebraic conjugates $\beta = \beta_1, \ldots, \beta_r$, and complex algebraic conjugates $\beta_{r+1}, \bar{\beta}_{r+1}, \ldots, \beta_{r+s}, \bar{\beta}_{r+s}$. The *conjugate map* $\Psi : \mathbb{Q}(\beta) \to \mathbb{R}^{r-1} \times \mathbb{C}^s \simeq \mathbb{R}^{d-1}$ is defined by $x \mapsto (x^{(2)}, \ldots, x^{(r)}, x^{(r+1)}, \ldots, x^{(r+s)})$ where $x^{(i)}$ denotes the image of x under the embedding of $\mathbb{Q}(\beta)$ into $\mathbb{C}$ that takes β to β_i. Hereafter, we identify $\mathbb{R}^{r-1} \times \mathbb{C}^s$ with $\mathbb{R}^{d-1}$.

Beta-tiles. Given a fractional part x_F of an element $x \in X_\beta$, and given $\cdots b_{-2}b_{-1}b_0$ in the past set $\mathcal{P}(x_F)$ (see Section 2.4), we define $\tilde{\Psi}(\cdots b_{-2}b_{-1}b_0) \in \mathbb{R}^{d-1}$ by $\tilde{\Psi}(\cdots b_{-2}b_{-1}b_0) := \sum_{i=0}^{\infty} b_{-i}\Psi(\beta^i)$. The series converges since the conjugates of β have absolute value less than 1 and the *prototiles* $\tilde{\Psi}(\mathcal{P}(x_F))$ correspond to the Rauzy pieces. More generally, versions of the Rauzy fractal and arithmetical coding when the Pisot number β is not a unit require either the inclusion of a non-archimedean component in the conjugate map (in the arithmetical setting) or the use of inverse limits (in the topological setting). See [109, 35, 21, 10] for more detail.

Non-periodic multi-tilings. Beta-tiles are analogues of Rauzy pieces. As in the symbolic substitution case, there is a natural multi-tiling that can be formed with the beta-tiles. The difference here is that this multi-tiling is not periodic. In fact, it is strongly connected with the dual multi-tiling of E^s discussed above for symbolic substitutions and as highlighted in [35, 36].

For $x_F = a_1 a_2 \cdots$ finite (that is, $a_i = 0$ for all i bigger than some k), the tile

$$\mathcal{T}(x_F) := \tilde{\Psi}(\mathcal{P}(x_F)) + \Psi(\pi(x_F))$$

is a translate of a prototile. (We recall that π is defined in (2.6).) Let points $\ell_i \in [0,1]$ be determined by $\{0 = \ell_0 < \ell_1 < \cdots < \ell_{m+\ell=1}\} = \{0\} \cup \{\pi(s^n(\mathbf{d}_\beta(\mathbf{1}-)))\}$. Then, if $x, y \in X_\beta$ have finite fractional parts and if $\pi(x_F), \pi(y_F) \in [\ell_{i-1}, \ell_i)$, we have $\mathcal{P}(x_F) = \mathcal{P}(y_F)$ so that $\mathcal{T}(y_F) = \mathcal{T}(x_F) + \Psi(\pi(y_F) - \pi(x_F))$ and thus there are (at most) $m + \ell$ different tiles up to translation.

The linear isomorphism of $\mathbb{Q}(\beta)$ given by multiplication by $1/\beta$ extends, after embedding $\mathbb{Q}(\beta)$ in $\mathbb{R}^{d-1}$ via Ψ, to an expanding linear isomorphism G of $\mathbb{R}^{d-1}$ and the tiles $\mathcal{T}(x_F)$ satisfy the substitution rule:

$$G(\mathcal{T}(x_F))) = \bigcup_{ax_F \text{ is admissible}} \mathcal{T}(ax_F).$$

The collection $\{\mathcal{T}(x_F) : x_F \text{ finite}\}$ is a multi-tiling of $\mathbb{R}^{d-1}$ invariant under this substitution rule ([3]). We will see in Section 6 equivalent arithmetic formulations of the following result.

Theorem 3.7 ([3, 21, 6]). *Let β be a Pisot unit of degree d. If the collection $\{\mathcal{T}(x_F) : x_F \text{ finite}\}$ is a tiling of $\mathbb{R}^{d-1}$, then $(\Omega_{\phi_\beta}, \mathbb{R})$ has pure discrete spectrum.*

Markov partitions. A Markov partition for a hyperbolic toral automorphism provides an explicit measurable conjugacy with a subshift of finite type by recording the itinerary of a point under the action of the automorphism with respect to the partition. According to [44], all hyperbolic toral automorphisms admit Markov partitions, but they can be rather difficult to find.

Sidorov describes in [105] the following procedure for constructing Markov partitions for certain total automorphisms (see also [93] and [103, 104]). In [70] and [106] this construction arises in the general framework of arithmetic dynamics and arithmetic codings. Given a Pisot unit β of degree d, let M be the companion matrix of its minimal polynomial. There is then an associated hyperbolic automorphism, F_M, of the d-torus $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ defined by $F_M(x + \mathbb{Z}^d) := Mx + \mathbb{Z}^d$.

With E^s and E^u denoting the stable and unstable spaces of M, the stable and unstable manifolds of $0 \in \mathbb{T}^d$ under F_M are $W^s(0) := E^s + \mathbb{Z}^d$ and $W^u(0) := E^u + \mathbb{Z}^d$, and the *homoclinic group* of 0 is the intersection $\mathcal{H} = W^s(0) \cap W^u(0)$. A point $\bar{y} \in \mathcal{H}$ is a *fundamental homoclinic point* if its orbit $\{F_M^n(\bar{y}) : n \in \mathbb{Z}\}$ generates the group $\mathcal{H}$. Given any $\bar{y} \in \mathcal{H}$, let $h_{\bar{y}} : X_\beta \to \mathbb{T}^d$ be defined by

$$h_{\bar{y}}(\cdots x_{-1}x_0 \bullet x_1 \cdots) := \sum_{n\in\mathbb{Z}} x_n F_M^{-n}(\bar{y}).$$

The map $h_{\bar{y}}$, called an *arithmetical coding* of F_M, is continuous, a.e. m-to-1 for some $m \in \mathbb{N}$, and semi-conjugates the shift s on X_β with F_M. Moreover, if $\underline{x}, \underline{y}$, and $\underline{x+y}$ denote the (greedy) β-expansions of the non-negative real numbers x, y and $x + y$, then $h_{\bar{y}}(\underline{x+y}) = h_{\bar{y}}(\underline{x}) + h_{\bar{y}}(\underline{y})$.

For each past set $\mathcal{P}$, let

$$[\mathcal{P}] := h_{\bar{y}}(\{x \in X_\beta : x_P = \mathcal{P}\}).$$

The intersections of $E^s \simeq \mathbb{R}^{d-1}$ with the sets $[\mathcal{P}]$ under the immersion which maps $x \in E^s$ to $x + \mathbb{Z}^d \in W^s(0)$ determine a degree m multi-tiling. If $m = 1$, the sets $[\mathcal{P}]$, where $\mathcal{P}$ are past sets constructed from a fundamental homoclinic point $\bar{y}$, form a Markov partition for the hyperbolic automorphism F_M associated with M ([105]). When the associated beta-substitution ϕ_β is irreducible, the sets $[\mathcal{P}]$ actually match with the sets $\mathcal{C}_i$ of the periodic multi-tiling of $\mathbb{R}^d$ built from the Rauzy pieces $\mathcal{R}_i$ introduced in Section 3.4. One similarly determines Markov partitions for toral automorphisms provided by the substitution matrix of unimodular irreducible Pisot substitutions when the tiling system has pure discrete spectrum; see, *e.g.*, [95, 65, 15].

It is proved in [105] that $m = 1$ (for a fundamental homoclinic point $\bar{y}$) if and only if $(\Omega_{\phi_\beta}, \mathbb{R})$ has pure discrete spectrum ([3, 21]). This property is itself equivalent to β satisfying the Property (W) discussed in Section 6 (see also [93] in the same vein).

4. The Pisot Substitution Conjecture

In this section, we focus on the one-dimensional case $n = 1$, that is, we consider either symbolic substitutions, beta-numeration, or tile substitutions defined in $\mathbb{R}$.

Recall that an algebraic integer $\lambda > 1$ is a *Pisot number* if all its algebraic conjugates α other than λ itself satisfy $|\alpha| < 1$. In terms of its influence on spectral properties of associated dynamical systems, the key property of a Pisot number λ is that the distance from λ^n to the nearest integer tends to zero as n tends to infinity. (Conversely, if λ is any algebraic number bigger than 1 with this property, then λ must be a Pisot number ([92]); it is a conjecture of Pisot that no transcendental number has this property.) A primitive substitution ϕ is said to be *Pisot* if its Perron–Frobenius eigenvalue λ is a Pisot number. By *Pisot dynamics* we mean, loosely, the shift dynamics on the symbolic space associated with a Pisot substitution, or the translation dynamics on a substitution tiling space with expansion $\Lambda = (\lambda)$, λ a Pisot number (or, more generally, Λ with the Pisot property, see Section 7), etc.

For the shift, or translation, dynamics associated with a substitution to have pure discrete spectrum a Pisot condition is necessary. The Pisot Substitution Conjectures have grown out of attempts to answer the question: what additional conditions guarantee that Pisot dynamics have pure discrete spectrum? In previous sections we have considered several types of Pisot dynamical systems, all arising from substitutions in a symbolic or geometrical context. The most easily recognized feature shared by all known examples that fail to have pure discrete spectrum is reducibility of the characteristic polynomial of the substitution matrix. Correspondingly, the most basic of the Pisot Substitution Conjectures, formulated separately for the symbolic and tiling contexts, are the following.

Conjecture 4.1 (Pisot Substitution Conjecture: symbolic substitutive case). If ϕ is an irreducible Pisot substitution then the substitutive system (X_ϕ, s) has pure discrete spectrum.

Conjecture 4.2 (Pisot Substitution Conjecture: tiling of the line case). If ϕ is an irreducible Pisot substitution then the tiling dynamical system $(\Omega_\phi, \mathbb{R})$ has pure discrete spectrum.

That these two conjectures are equivalent is a consequence of a theorem of Clark and Sadun ([45]). Indeed, (X_ϕ, s) has pure discrete spectrum if and only if the $\mathbb{R}$-action on the suspension $\hat{X}_\phi$, with constant roof function, has pure discrete spectrum. If the constant value c of the roof function is chosen so that $c(1, \ldots, 1)-\ell$, ℓ the positive left eigenvector of the substitution matrix M_ϕ giving the tile lengths, is in the left-contracting space of the substitution matrix M_ϕ (so that $(c(1, \ldots, 1)-\ell)M_\phi^m \to 0$ as $m \to \infty$) then, according to Theorem 3.1 of [45], the $\mathbb{R}$-actions on $\hat{X}_\phi$ and Ω_ϕ are topologically conjugate. Such a choice of c is always possible in the irreducible case as then the left contracting space has codimension one.

The Thue–Morse substitution ($1 \mapsto 12$, $2 \mapsto 21$) provides a simple example showing that the above conjectures are false if the hypothesis of irreducibility is

dropped. No such examples have been found for β-substitutions (see Section 2.4) and in this setting there is a stronger conjecture.

Conjecture 4.3 (Pisot Substitution Conjecture: β-substitution case). If β is a Pisot number with associated β-substitution ϕ_β then the tiling dynamical system $(\Omega_{\phi_\beta}, \mathbb{R})$ has pure discrete spectrum.

For counterexamples to the above (with reducible substitution matrix M_{ϕ_β}) for the substitutive system (X_{ϕ_β}, s), rather than the tiling dynamics, see [49]. If the tiling dynamical system $(\Omega_{\phi_\beta}, \mathbb{R})$ has pure discrete spectrum, these examples show that the symbolic substitutive system (X_{ϕ_β}, s) may not. Nonetheless, in this situation (X_{ϕ_β}, s) is at least measurably isomorphic with an induced system of a group translation (see [21]).

The assumption of 'irreducibility' in Conjectures 4.1 and 4.2 is unnatural in the following sense. It is easy to take an irreducible Pisot substitution and 'rewrite' it to obtain another substitution that is not irreducible but has topologically conjugate dynamics ([24]). One substitution will satisfy the hypotheses while the other won't, yet their spectral properties are identical. In fact, if two one-dimensional substitution tiling spaces are homeomorphic, then the corresponding $\mathbb{R}$-actions are (up to rescaling) conjugate ([31], and [71] for higher dimensions). It is thus desirable to find a topological condition to replace the assumption of irreducibility of the substitution matrix.

The substitution homeomorphism $\Phi : \Omega_\phi \to \Omega_\phi$ induces a linear isomorphism $\Phi^* : H^1(\Omega_\phi) \to H^1(\Omega_\phi)$ on the (Čech, with rational coefficients) cohomology of Ω_ϕ. The expansion λ is an eigenvalue of this isomorphism so the dimension of $H^1(\Omega)$ is at least the algebraic degree d of λ ([26]). Thus Φ^* is irreducible (that is, its characteristic polynomial is irreducible over $\mathbb{Q}$) if and only if the dimension of $H^1(\Omega_\phi)$ equals d. If Φ^* is irreducible, we will say that ϕ is *homologically irreducible* and that ϕ is a *homological Pisot substitution* if ϕ is also a Pisot substitution.

Conjecture 4.4 (Homological Pisot Substitution Conjecture). If ϕ is a homological Pisot substitution whose expansion is an algebraic unit then the tiling dynamical system $(\Omega_\phi, \mathbb{R})$ has pure discrete spectrum.

By [46], and as in the comments above following Conjecture 4.2, if ϕ is homological Pisot, then (X_ϕ, s) has pure discrete spectrum if and only if $(\Omega_\phi, \mathbb{R})$ does: a similar remark applies to the Coincidence Rank Conjecture below.

Conjecture 4.4 is neither weaker nor stronger than Conjectures 4.1 and 4.2 (with an additional assumption of unimodularity). However, the only way an irreducible substitution ϕ can fail to be homologically irreducible is for there to be an 'asymptotic cycle' of arc components in Ω_ϕ which is associated with a root of unity eigenvalue for Φ^* ([25]).

There are counterexamples to Conjecture 4.4 if the expansion is not assumed to be a unit ([33]) but the conjecture can be extended in the following way. The *norm* of an algebraic integer λ is the product of λ with all its algebraic conjugates. So λ is a unit if and only if its norm is ± 1 and if λ is the PF-eigenvalue of a

substitution matrix with irreducible characteristic polynomial M_ϕ, then the norm of λ is $\pm \det(M_\phi)$.

See Section 3.1 for the definition of the 'coincidence rank' of a substitution and recall that the coincidence rank of a Pisot substitution equals one if and only if the system $(\Omega_\phi, \mathbb{R})$ has pure discrete spectrum. Thus the following conjecture extends the Homological Pisot Substitution Conjecture to the non-unit case.

Conjecture 4.5 (Coincidence Rank Conjecture). If ϕ is a homological Pisot substitution with expansion λ then the coincidence rank of ϕ divides the norm of λ.

5. Coincidences

We present here several variations of the notion of coincidence. Strong coincidences (see Section 5.1), geometric coincidences (Section 5.2) and balanced pairs (Section 5.3) are defined in the one-dimensional case $n = 1$, whereas coincidences in higher dimensions are handled in Section 5.4 with the notions of modular coincidence, overlap coincidence and algebraic coincidence. We then revisit the notion of algebraic coincidence in the particular case of beta-numeration in Section 5.5.

5.1. Strong coincidences

A primitive substitution ϕ is said to satisfy the *strong coincidence condition* if each pair of ϕ-periodic points $x = (\cdots x_{-1}x_0x_1\cdots), y = (\cdots y_{-1}y_0y_1\cdots) \in X_\phi$ are *strongly coincident*, *i.e.*, there is $n \in \mathbb{N}$ so that $x_n = y_n$ and the abelianizations $f(x_0\cdots x_n)$ and $f(y_0\ldots y_n)$ are equal. This combinatorial condition, originally due to Arnoux and Ito [14], is an extension of a similar condition considered by Kamae [67] and Dekking [47] in the case of *constant length substitutions* (that is, when $|\phi(a)| = |\phi(b)|$ for all $a, b \in \mathcal{A}$). In [47], Dekking proves that a constant length substitution having trivial height satisfies the strong coincidence condition if and only if (X_ϕ, s) has pure discrete spectrum.

Conjecture 5.1 (Strong coincidence Conjecture). Every irreducible Pisot substitution satisfies the strong coincidence condition.

It is not known whether or not the Coincidence Conjecture is equivalent to the Pisot Substitution Conjecture (that is, if strong coincidence implies pure discrete spectrum). We formulate now a geometrical version of coincidence that is equivalent to pure discrete spectrum.

5.2. Geometric coincidences

The idea of associating a geometrical strand with an element of a symbolic substitution space originates with Arnoux and Itô in [14]. There, the authors present a geometrical version of a substitution: letters become segments in $\mathbb{R}^d$ with integer vertices, and their images under substitution become 'broken lines'. This allows for a dualization of the substitution: dual to a segment is a $d-1$ cell which, under the dual substitution becomes a piece of a broken hyperplane. With these tools, the authors prove that if an irreducible unimodular Pisot substitution satisfies

the strong coincidence condition, then its substitutive system is measurably conjugate to a domain exchange in $\mathbb{R}^{d-1}$ and semi-conjugate to a rotation on the $(d-1)$-torus. The idea is further developed in [30, 21], where the segments are no longer required to have vertices on the integer lattice. With this innovation, the tiling space itself, in the form of strand space, emerges as the global attractor of a geometrical substitution. Each element of the tiling space now corresponds uniquely to a strand in $\mathbb{R}^d$ and the map taking the tiling to a vertex of its strand, modulo $\mathbb{Z}^d$, semi-conjugates the tiling flow with a Kronecker action on $\mathbb{T}^d$ and the substitution homeomorphism with a hyperbolic automorphism of $\mathbb{T}^d$.

For simplicity in what follows, we assume that ϕ is an irreducible and unimodular Pisot substitution on d letters; also, we construct strand space as an orbit closure, rather than as a global attractor.

To begin with, for $x \in \mathbb{R}^d$ and e_i a standard unit vector, we call the set

$$[x, i] := \{x + te_i : 0 \le t \le 1\}$$

a *segment of type* i with *initial vertex* x and *terminal vertex* $x + e_i$. Suppose that $u \in X_\phi$ is ϕ-periodic. Employing the abelianization map f, the *strand associated with* u is the union of segments:

$$S_u = \bigcup_{N\in\mathbb{N}} [f(u_0 \cdots u_{N-1}), u_N] \cup \bigcup_{N\in\mathbb{N}} [-f(u_{-N-1} \cdots u_{-1}), u_{-N-1}] .$$

The tiling $\mathcal{T}_u$ is recovered from the strand S_u by projection onto the unstable space $E^u = \{tr : t \in \mathbb{R}\}$ spanned by the right PF-eigenvector r of the abelianization M_ϕ of ϕ [34].

The *strand space* associated with ϕ is the translation-orbit closure

$$\Sigma_\phi := \overline{\{S_u - tr : t \in \mathbb{R}\}},$$

the closure being taken in the topology in which two strands are close if their intersections with a large closed ball about the origin are Hausdorff-close. Given a segment $[x, i]$, let

$$\Phi([x, i]) := \cup_{j=1,\ldots,k}[M_\phi x + f(i_1 \cdots i_{j-1}), i_j],$$

where $\phi(i) = i_1 \cdots i_k$. Then, for $S = \cup_j [x_j, l_j] \in \Sigma_\phi$, $\Phi(S) := \cup_j \Phi([x_j, l_j])$ defines a self-homeomorphism of Σ_ϕ. Strand space is simply another presentation of tiling space: the $\mathbb{R}$-action $S \mapsto S - tr$ and the homeomorphism Φ are conjugated with the corresponding dynamics on Ω_ϕ through projection of strands to tilings of E^u. Strand space has two advantages: there is an easily defined *geometric realization*, $g : \Sigma_\phi \to \mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ by $g(S) = x + \mathbb{Z}^d$, where x is a vertex of a segment in S; and there is a *geometric coincidence condition*, by means of which pure discrete spectrum can be checked. To formulate the latter, we say that segments $[x, i]$ and $[y, j]$ are *coincident* if there is $n \in \mathbb{N}$ so that $\Phi^n([x, i]) \cap \Phi^n([y, j])$ contains a segment. The substitution ϕ *satisfies the geometric coincidence condition* if for all $i \in \mathcal{A}$ and for all $x \in \mathbb{Z}^d$ and $j \in \mathcal{A}$ so that the interior of $[x, j]$ meets the stable space, E^s, of M_ϕ, it happens that $[0, i]$ and $[x, j]$ are coincident. This condition is called *super coincidence* in [65]. The equivalence of pure discrete spectrum

and geometric/super coincidence for irreducible unimodular Pisot substitutions appears in [30] and [65] and is generalized to the reducible and non-unimodular setting in [21] and [50].

Theorem 5.2 (Geometric coincidence condition and geometric realization). *If ϕ is an irreducible unimodular Pisot substitution then the geometric realization $g : \Sigma_\phi \to \mathbb{T}^d$ semi-conjugates Φ with the automorphism of $\mathbb{T}^d$ induced by M_ϕ and the $\mathbb{R}$-action on Σ_ϕ with a Kronecker action on $\mathbb{T}^d$. The following are equivalent:*

(i) *$(\Omega_\phi, \mathbb{R})$ has pure discrete spectrum,*
(ii) *g is almost everywhere one-to-one,*
(iii) *ϕ satisfies the geometric coincidence condition.*

The Kronecker action in Theorem 5.2 is the maximal equicontinuous factor of the $\mathbb{R}$-action on Ω_ϕ mentioned in Section 3.1; more will be said about this in Section 7.1. Equivalence of the geometric coincidence condition with measurable conjugacy of the tiling and Kronecker dynamical systems can be found in [30] and [21]. The map g in Theorem 5.2 is, in general, almost everywhere m-to-one with m equal to the maximal cardinality of a collection of segments $\{[x_i, j_i]\}$ having the properties: each of these segments meets E^s in its interior; $x_i - x_k \in \mathbb{Z}^d$ for all i, k; and $[x_i, j_i]$ and $[x_k, j_k]$ are not coincident for $i \neq k$. This m is the coincidence rank of ϕ (see Section 3.1 for the definition and Conjecture 4.5). The geometric and super coincidence conditions are versions of the more general *overlap coincidence* condition introduced in [113] for tilings of the plane and specialized to one-dimensional tilings in [111]. Akiyama and Lee have automated an overlap coincidence algorithm in [6]. According to the formalism of dual substitutions developed in [14], dual versions of these conditions have been expressed in [14, 65, 108, 39].

5.3. The balanced pair algorithm

The following *balanced pair algorithm* is a purely combinatorial adaptation of overlap coincidence presented in [111] for application to Pisot substitutive systems.

The balanced pair algorithm for establishing pure discrete spectrum originated with [87, 83, 96]. A pair of finite words $(x, y) \in \mathcal{A}^* \times \mathcal{A}^*$ is said to be a *balanced pair* if the abelianizations $f(x)$ and $f(y)$ are equal. A *coincidence* is a one-letter balanced pair (a, a). A balanced pair is *irreducible* if it cannot be properly factored as a product (with respect to the multiplication $(x, y)(u, v) := (xu, yv)$) of balanced pairs. Clearly, each balanced pair can be factored uniquely as a product of irreducible balanced pairs. Suppose that $u = vwv\cdots$ is a fixed (or periodic) word for the substitution ϕ. One says that *the balanced pair algorithm for ϕ terminates with coincidence* if only finitely many distinct irreducible factors occur in factorizations of all balanced pairs of the form $(\phi^n(vw), \phi^n(wv))$, $n \in \mathbb{N}$, and for each such occurring irreducible factor (x, y), $(\phi^k(x), \phi^k(y))$ has a coincidence in its irreducible factorization for some $k \in \mathbb{N}$. For more on the balanced pair algorithm, and the theorem below, see [111].

Theorem 5.3 (Balanced pair algorithm). *Given an irreducible Pisot substitution ϕ, the substitutive dynamical system (X_ϕ, s) has pure discrete spectrum if and only if the balanced pair algorithm for ϕ terminates with coincidence.*

In fact, it suffices to check termination with coincidence of the balanced pair algorithm starting from any particular seed of the form (ij, ji), $i \neq j \in \mathcal{A}$. For this, and extensions to the reducible setting, as well as geometric versions for tilings in arbitrary dimension, see [86] and [32].

5.4. Coincidences in higher dimensions

There are various notions of coincidence for n-dimensional substitution tilings which characterize pure discrete spectrum of the tiling dynamical system. We mention a few of these here.

Modular coincidence. In the case that the underlying structure of a substitution tiling is a lattice (this is the analog of a constant length substitution tiling in one-dimension), it is easy to check whether or not the tiling system has pure discrete spectrum by checking for 'modular coincidence', as we explain now.

Suppose that $\underline{\varGamma} = \bigcup_{i=1}^m \varGamma_i \times \{i\}$ is a primitive substitution Delone multi-color set with expansion map Λ and that $L = \bigcup_{1\leq i\leq m} \varGamma_i$ is a lattice (that is, a co-compact discrete subgroup of $\mathbb{R}^n$). Let $L_i := \langle \varGamma_i - \varGamma_i \rangle$ be the Abelian group generated by $\varGamma_i - \varGamma_i$ and let $L' := L_1 + L_2 + \cdots + L_m$. For $a \in L$, Φ_{ij} as in (2.4), and $f \in \Phi_{ij}$, let $t(f) \in \mathcal{D}_{ij}$ be so that $f(x, j) = (\Lambda x + t(f), i)$. Set

$$\begin{aligned} \Phi_{ij}[a] := \{f \in \Phi_{ij} :\ & \Lambda y + t(f) \equiv a \bmod \Lambda L', \\ & \text{where } \varGamma_j \subset y + L'\} \\ = \{f \in \Phi_{ij} : \operatorname{supp}(\varGamma_j \times \{j\}) & \subset a + \Lambda L')\}. \end{aligned} \tag{5.1}$$

Then

$$\bigcup_{i,j\leq m} \bigcup_{f\in\Phi_{ij}[a]} \operatorname{supp}(\varGamma_j \times \{j\}) = a + \Lambda L'.$$

Let $\Phi[a] := \cup_{1\leq i,j\leq m} \Phi_{ij}[a]$.

We say that $\underline{\varGamma}$ admits *a modular coincidence relative to* $\Lambda L'$ if $\Phi[a]$ is contained entirely in one row of Φ for some $a \in L$. It is easy to see that $\underline{\varGamma}$ admits a modular coincidence relative to $\Lambda L'$ if and only if $(a + \Lambda L') \subset \varGamma_i$ for some $1 \leq i \leq m$.

Theorem 5.4 ([76]). *Let $\mathcal{T}$ be a primitive substitution tiling in $\mathbb{R}^n$ with associated Delone multi-color set $\underline{\varGamma}_\mathcal{T} = \bigcup_{i=1}^m (\varGamma_i \times \{i\})$ and suppose that $L = \bigcup_{1\leq i\leq m} \varGamma_i$ is a lattice. Let $L' = L_1 + \cdots + L_m$, where $L_i = \langle \varGamma_i - \varGamma_i \rangle$. The following are equivalent:*

(i) *$(\Omega_\mathcal{T}, \mathbb{R}^n)$ has pure discrete spectrum.*
(ii) *A modular coincidence relative to $\Lambda^M L'$ occurs in Φ^M for some M.*
(iii) *Each $\varGamma_i$ is a regular model set for $1 \leq i \leq m$.*

The upper bound for the number of iterations to check modular coincidence is given in [56] where the underlying structure is on a lattice and in [6] for more general substitution tilings.

Algebraic coincidences. For substitution tilings in $\mathbb{R}^n$ whose associated Delone sets may not be on lattices, the notion of algebraic coincidence (defined below) generalizes the notion of modular coincidence and provides an opportunity for a concise expression of the connection between model sets and pure discrete spectrum for substitution tiling spaces.

Let $\underline{\varGamma}$ be a primitive substitution Delone multi-color set with an expansive map Λ. Let $\Xi(\underline{\varGamma}) = \bigcup_{j\le m}(\varGamma_j - \varGamma_j)$. We say that $\underline{\varGamma}$ admits an *algebraic coincidence* if there exist $M \in \mathbb{Z}_+$ and $\xi \in \varGamma_i$ for some $i \le m$ such that $\xi + \Lambda^M \Xi(\underline{\varGamma}) \subset \varGamma_i$.

Theorem 5.5 ([75]). *Let $\mathcal{T}$ be a primitive substitution tiling with FLC and let $\underline{\varGamma}_{\mathcal{T}} = \bigcup_{i=1}^m (\varGamma_i \times \{i\})$ be an associated substitution Delone multi-color set. Then the following are equivalent:*

(1) *$(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ has pure discrete spectrum.*
(2) *$\underline{\varGamma}_{\mathcal{T}}$ admits an algebraic coincidence.*
(3) *Each $\varGamma_i$ is an inter-model set.*

Overlap coincidences. Unlike modular coincidence, algebraic coincidence (for substitution tilings whose associated Delone sets are not on lattices) is not so easily checked. For the computation of pure discrete spectrum, the condition of 'overlap coincidence', which we define now, proves to be more convenient.

Let $\mathcal{T}$ be a tiling and let

$$\Xi(\mathcal{T}) := \{v \in \mathbb{R}^n : \text{ there is } T \in \mathcal{T} \text{ with } T + v \in \mathcal{T}\} \tag{5.2}$$

be the set of *return vectors* for $\mathcal{T}$. A triple (T, y, S), with $T, S \in \mathcal{T}$ and $y \in \Xi(\mathcal{T})$, is called an *overlap* if $\operatorname{supp}(y + T) \cap \operatorname{supp}(S)$ has non-empty interior.

An overlap (T, y, S) is a *coincidence* if $y + T = S$. The *support* of an overlap (T, y, S) is $\operatorname{supp}(T, y, S) = \operatorname{supp}(y+T) \cap \operatorname{supp}(S)$. Let $\mathcal{O} = (T, y, S)$ be an overlap. Recall that for a tile-substitution Φ, $\Phi(y + T) = \Lambda y + \Phi(T)$ is a patch of $\Lambda y + \mathcal{T}$, and $\Phi(S)$ is a $\mathcal{T}$-patch; moreover,

$$\operatorname{supp}(\Lambda y + \Phi(T)) \cap \operatorname{supp}(\Phi(S)) = \Lambda(\operatorname{supp}(T, y, S)).$$

For each $l \in \mathbb{Z}_+$,

$$\Lambda^l(\mathcal{O}) = \{(T', \Lambda^l y, S') : T' \in \Phi^l(T), S' \in \Phi^l(S), \operatorname{supp}(\Lambda^l y + T') \cap \operatorname{supp}(S') \neq \emptyset\}.$$

We say that a substitution tiling $\mathcal{T}$ admits an *overlap coincidence* if there exists $l \in \mathbb{Z}_+$ such that for each overlap $\mathcal{O}$ in $\mathcal{T}$, $\Lambda^l(\mathcal{O})$ contains a coincidence. We recall that the Meyer property was introduced in Section 3.2.

Theorem 5.6 ([81, Thm. 4.7 and Lemma A.9]). *Let $\mathcal{T}$ be a primitive substitution tiling which has the Meyer property. Then $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ has pure discrete spectrum if and only if $\mathcal{T}$ admits an overlap coincidence.*

When a tiling has the Meyer property, the number of equivalence classes of overlaps is finite. Thus, once all equivalences of overlaps are found, by applying the substitution to each overlap it can be determined if the overlap coincidence condition holds. When the dimension of the tiling is more than 2, however, it is not easy to check if a given triple $\mathcal{O} = (T, y, S)$ is an overlap. Instead, one can consider potential overlaps (T', x, S') for which $T' + x$ and S' are within certain distance. Then, computing how many potential overlaps come from each potential overlap after substitution, one can tell which potential overlaps are actually overlaps by simple computation of spectral radii. This easy procedure is justified by proving that the (slightly modified) Hausdorff dimension of the tile boundary of the self-affine tiling is strictly less than the dimension of the space. This leads to the algorithm of Akiyama and Lee ([6]) for determining overlap coincidence. We recall that $\Xi(\mathcal{T})$ is the set of return vectors for the tiling $\mathcal{T}$ such as defined in (5.2).

Theorem 5.7. *Let $\mathcal{T}$ be a primitive substitution tiling for which $\Xi(\mathcal{T})$ is a Meyer set and the digit sets of the tile substitution are provided. Then there is a terminating algorithm determining overlap coincidence.*

5.5. Beta-numeration: Property (W) and algebraic coincidences

We now return to the one-dimensional case in the particular framework of beta-numeration with the introduction of the condition called Property (W), which, while not directly stated in terms of coincidences, is nonetheless closely related to the notion of algebraic coincidence introduced in the previous section.

The β-expansion $\mathbf{d}_\beta(x)$ is *finite*, if there is n_0 such that $d_n(x) = 0$ for $n > n_0$. The image by π (defined in (2.6)) of such an element is written as $d_{-m}(x)d_{-m+1}(x)\cdots d_{n_0}(x)$ for simplicity. Let $\mathrm{Fin}(\beta)$ denote the set of all $x \geq 0$ whose β-expansion is finite. We consider several properties concerning $\mathrm{Fin}(\beta)$.

- (W) For each $x \in \mathbb{Z}[1/\beta] \cap [0, \infty)$ and each $\varepsilon > 0$, there are $y, z \in \mathrm{Fin}(\beta)$ with $z < \varepsilon$ and $x = y - z$.
- (H) $\mathbb{Z}[1/\beta] \cap [0, 1) \subset (\mathrm{Fin}(\beta) \cap [0, \beta)) - (\mathrm{Fin}(\beta) \cap [0, 1))$.

The conditions (W) and (H) are equivalent (see [9]). The notation (H) derives from Hollander [63], who showed that the slightly stronger property $\mathbb{Z}[1/\beta] \cap [0, 1) = (\mathrm{Fin}(\beta) \cap [0, 1)) - (\mathrm{Fin}(\beta) \cap [0, 1))$ implies that (X_{ϕ_β}, s) has pure discrete spectrum when X_β is a subshift of finite type and ϕ_β is irreducible.

Recall from Section 3.5 that if β is a Pisot unit of degree d then $\{\mathcal{T}(x_F) : x_F \text{ finite}\}$ is a multi-tiling of $\mathbb{R}^{d-1}$. The system $(\Omega_{\phi_\beta}, \mathbb{R})$ has pure discrete spectrum if and only if the degree (multiplicity) of this multi-tiling is one. If ϕ_β is irreducible, that is, if $d = n$, this is equivalent to saying that the system (X_{ϕ_β}, s) has pure discrete spectrum.

We now introduce a topological criterion:

- (Ex) There is an exclusive inner point in $\mathcal{T}(x_F)$ for some finite fractional part x_F (*i.e.*, $\pi(x_F) \in \mathbb{Z}[1/\beta] \cap [0, 1)$) of an $x \in X_\beta$.

Here an exclusive inner point means it does not lie in any tile $\mathcal{T}(y_F)$ with $y_F \neq x_F$. It is shown in [3] that the degree of the multi-tiling is one if and only if (Ex) holds,

and that (W) and (Ex) are equivalent, that is, (W), (Ex), and pure discreteness of $(\Omega_{\phi_\beta}, \mathbb{R})$ are equivalent.

Indeed, let us call an element of $\mathcal{P}(x_F)$ *finite* if it has the form $0^\infty y$. The map π extends to such finite elements and $\Gamma_{x_F} = \{\pi(y) \mid 0^\infty y \in \mathcal{P}(x_F)\}$ is a set of control points for the tiling of the half-line obtained by iterating the tile substitution (2.7) on the tile $\pi(\mathcal{F}(x_I))$ with $I = \cdots 000$. Let $\mathbf{P} = \{x \in \mathbb{Z}[\beta] \cap [0,1) \mid (T_\beta^n(x))_{n=0,1,\ldots}$ is purely periodic$\}$. When β is a Pisot number, we easily see that $\mathbf{P}$ is a finite set and for each $x \in \mathbb{Z}[1/\beta] \cap [0,\infty)$ there is $k \in \mathbb{Z}$ such that $T_\beta^k(x) \in \mathbf{P}$. In fact, there is a uniform k so that every non-negative element z in $\Gamma_{x_F} - \Gamma_{x_F}$ satisfies $T_\beta^k(z) \in \mathbf{P}$. From this we are able to see the connection between (Ex) and the algebraic coincidence formulated by Lee [75] (and discussed in Section 5.4) as follows. By Proposition 1 of [3], (Ex) is equivalent to the existence of an element $x \in \mathbb{Z}[\beta]$ and a constant K_0 so that $\pi(x) \in \mathcal{P}(0^\infty)$ and $\pi(\beta^K u + x) \in \mathcal{P}(0^\infty)$ for any integer $K \geq K_0$ and any $u \in \mathbf{P}$. One readily sees that this condition is a special form of algebraic coincidence of the one-dimensional tiling. Thus Property (W) is equivalent to pure discreteness of the translation action on the hull of a one-dimensional tiling generated by β-expansion.

On the other hand, once we have a Pisot dual tiling (see Section 3.5), *i.e.*, the degree of the dual multi-tiling is one, we can immediately show algebraic coincidence for such a $(d-1)$-dimensional tiling. We see that $\pi(\mathcal{F}(0^\infty)) = \mathbb{Z}[\beta] \cap [0,1)$ and $\Phi(\mathbb{Z}[\beta] \cap [0,1))$ is the union of control points of Pisot dual tilings. Taking algebraic conjugates into consideration, to have algebraic coincidence, we only need to find another constant K_1 and an element $x \in \mathbb{Z}[\beta] \cap [0,1)$ such that $x + \beta^{-K}(\mathbb{Z}[\beta] \cap [0,1) - \mathbb{Z}[\beta] \cap [0,1)) \subset \mathbb{Z}[\beta] \cap [0,1)$ for $K \geq K_1$. This turns out to be trivial from the Parry condition introduced in Section 2.4.

Summing up, the Pisot Substitution Conjecture for β-substitutions is equivalent to (W) which is equivalent to the associated dual multi-tiling having degree one. Once the multi-tiling has degree one, the associated $(d-1)$-dimensional tiling dynamical system is pure discrete, as is $(\Omega_{\phi_\beta}, \mathbb{R})$.

6. Partial results toward pure discrete spectrum

Substitution case. Conjectures 4.1 and 4.2 have been established for symbolic substitutions on two letters ([111]) and are known to hold for some families of substitutions, such as, *e.g.*, the Arnoux–Rauzy substitutions ([32, 38, 41, 23]) and the substitutions associated with the Brun and Jacobi-Perron continued fraction algorithms ([40, 64, 41, 23]). They have also been checked, mainly by the methods of Section 5.3, for many more-or-less randomly chosen substitutions and special cases of substitutions on three letters ([7]). Otherwise, they remain wide open.

The Coincidence Rank Conjecture 4.5 is verified for degree one (that is, when the expansion factor λ is an integer) in [33] and for all Pisot λ in the case that the coincidence rank is two in [22]. Thus there are no counterexamples to the Homological Pisot Substitution Conjecture with coincidence rank two and any

coincidence rank two counterexample to the Pisot Substitution Conjectures 4.1 and 4.2 must have an asymptotic cycle of arc components.

Beta-numeration case. The context of beta-substitutions, being narrower, has seen more progress. Here we list known sufficient conditions for the validity of the beta-substitution case of the Pisot Substitution Conjecture which make use only of algebraic conditions on the Pisot number β. Let $x^d - \sum_{i=0}^{d-1} k_i x^i$ be the minimal polynomial of β. We recall that Property (W), which is equivalent to pure discrete spectrum of the tiling system, has been introduced in Section 5.5.

- The condition $k_{d-1} > \sum_{i=0}^{d-2} |k_i|$ implies (W) (see [9]).
- Pisot units with $d \leq 3$ satisfy (W) (see [9]).
- The condition $\mathbf{d}_\beta(1-) = (c_1 c_2 \cdots c_m)^\infty$ implies (W) (see [23]).

Moreover, recall that $\mathrm{Fin}(\beta)$ denotes the set of all $x \geq 0$ whose β-expansion is finite. In addition to the (W) and (H) properties introduced in Section 5.5, we consider two stronger properties concerning $\mathrm{Fin}(\beta)$:

- (F) $\mathrm{Fin}(\beta) \supset \mathbb{Z}[1/\beta] \cap [0, \infty)$;
- (PF) $\mathrm{Fin}(\beta) \supset \mathbb{Z}_+[1/\beta]$ where $\mathbb{Z}_+ = \mathbb{Z} \cap [0, \infty)$.

The conditions (F) and (PF) imply (W) (see [9]). They were introduced by Frougny and Solomyak in [57] who proved the following:

- $c_1 \leq c_2 \leq c_3 \cdots$ implies (F) or (PF) ([57]).

  Those β with property (PF) but not (F) are characterized in [4].

The finiteness condition (F) means all possible candidates have finite β-expansion and it is equivalent to state $\mathrm{Fin}(\beta) = \mathbb{Z}[1/\beta] \cap [0, \infty)$. It is useful in many situations with ergodic and number theoretical flavor: Akiyama showed in [1] that every sufficiently small rational number has purely periodic T_β-orbit under (F) and, in [93], Praggastis constructed Markov partitions for toral automorphisms related to β-expansion under condition (F).

The property $\mathbb{Z}_+ \subset \mathrm{Fin}(\beta)$ implies that β is a Pisot number (Proposition 1 in [2]), thus a number β satisfying (F) is a Pisot number, but the converse is false. For example, if β has property (F) then it cannot have another positive conjugate. The characterization problem of Pisot numbers with property (F) is difficult and has been transformed into a problem of shift radix systems (see [12, 11]). The idea of a shift radix system is essentially due to Gilbert [59] and Hollander [63] (see [8]). For a version of the finiteness property (F) in the symbolic substitution case and its relation with topological properties of Rauzy fractals, see [35, 108, 39].

7. The Pisot property and hyperbolicity in higher dimensions

For the $\mathbb{R}$-action on a one-dimensional substitution tiling space to have pure discrete spectrum it is necessary that the expansion λ be a Pisot number ([113]). Generally, for the $\mathbb{R}^n$-action on an n-dimensional substitution tiling space to have pure discrete spectrum it is necessary that the total number (with multiplicity) of

algebraic conjugates η of eigenvalues of the linear expansion Λ with $|\eta| > 1$, equals n. The term *Pisot property* is designed to capture this condition.

7.1. Pisot families and the Pisot property

First let us remark that if Φ is an n-dimensional primitive tiling substitution with expansion Λ then the eigenvalues of Λ are all algebraic integers, so speaking of their algebraic conjugates makes sense (see [69] for the diagonalizable case and [72] for the general result). Let $J[\lambda, r]$ denote a real Jordan block with either real eigenvalue λ and size $r \times r$ or complex eigenvalues $\lambda, \bar{\lambda}$ and size $2r \times 2r$. Then Λ is said to have the *Perron property* if, whenever $J[\lambda, r]$ occurs in the real Jordan form of Λ with multiplicity k and λ' is an algebraic conjugate of λ with $|\lambda'| \geq |\lambda|$, then $J[\lambda', r]$ is also a block in the real Jordan form of Λ with multiplicity at least k (we consider $J[\lambda, r]$ and $J[\bar{\lambda}, r]$ to be the same). The terminology is due to Kwapisz who proves in [72], generalizing a result of Kenyon and Solomyak [69], that linear expansions for primitive tile substitutions must have the Perron property. Let us say that Λ has the *Pisot property* if, whenever $J[\lambda, r]$ is a block in the real Jordan form of Λ with multiplicity k and λ' is an algebraic conjugate of λ with $|\lambda'| \geq 1$, then $J[\lambda', r]$ also occurs in the real Jordan form of Λ and with multiplicity k.

If, in the definition of the Pisot property, one drops reference to Jordan blocks and speaks instead only of eigenvalues, the *Pisot family condition* results (we caution the reader that the definition of Pisot family is somewhat variable in the literature). But the Pisot family condition plus the Perron property is equivalent to the Pisot property, so in light of the Kwapisz result cited above, if Λ is the expansion for a primitive FLC tile substitution, then Λ has the Pisot property if and only if Λ satisfies the Pisot family condition.

Every primitive FLC substitution tiling dynamical system $(\Omega_\Phi, \mathbb{R}^n)$ has a maximal equicontinuous factor $(\hat{\mathbb{T}}, \mathbb{R}^n)$ with factor map $g : \Omega_\Phi \to \hat{\mathbb{T}}$ with g a.e. m-to-1 for some $m \in \mathbb{N} \cup \{\infty\}$, where $\hat{\mathbb{T}}$ is a torus or solenoid and the $\mathbb{R}^n$-action is a Kronecker action. The number $m = cr(\Phi)$ is the coincidence rank of Φ ([28]). Also, by [113, 29, 21], for one-dimensional primitive tile substitutions Φ, $cr(\Phi) < \infty$ if and only if the expansion λ for Φ is a Pisot number.

Conjecture 7.1. If Φ is a primitive FLC tile substitution then $cr(\Phi) < \infty$ if and only if Φ has the Pisot property.

7.2. Pisot families and discrete spectrum

We sketch below the sufficiency of a strong form of the Pisot property in Conjecture 7.1; it is reasonable to expect that the full conjecture will follow along the lines of [72].

Let us consider two additional conditions on the expansion Λ:

[A1] The expansion map Λ is diagonalizable over $\mathbb{C}$.

[A2] All eigenvalues of Λ are algebraic conjugates with the same multiplicity.

Let J be the multiplicity of each eigenvalue of Λ. After a linear change of coordinates, we may write

$$\Lambda = \begin{bmatrix} \Psi_1 & \cdots & O \\ \vdots & \ddots & \vdots \\ O & \cdots & \Psi_J \end{bmatrix}$$

where $\Psi_j = \Psi$ for any $1 \le j \le J$, Ψ is an $m \times m$ matrix, and O is the $m \times m$ zero matrix. For each $1 \le j \le J$, let

$$H_j = \{0\}^{(j-1)m} \times \mathbb{R}^m \times \{0\}^{n-jm}.$$

We define $\boldsymbol{\alpha}_j \in H_j$ such that for each $1 \le d \le n$,

$$(\boldsymbol{\alpha}_j)_d = \begin{cases} 1 & \text{if } (j-1)m+1 \le d \le jm; \\ 0 & \text{else}. \end{cases} \tag{7.1}$$

In the following theorem, which is key for the proof of Theorem 7.3, $\mathcal{C}$ is a set of control points for $\mathcal{T}$ (see Section 2.3).

Theorem 7.2 ([78]). *Let $\mathcal{T}$ be a primitive FLC substitution tiling of $\mathbb{R}^n$ with expansion Λ satisfying* [A1] *and* [A2]. *Then there exists a linear isomorphism $\rho : \mathbb{R}^n \to \mathbb{R}^n$ such that*

$$\rho\Lambda = \Lambda\rho \quad \text{and} \quad \mathcal{C} \subset \rho(\mathbb{Z}[\Lambda]\boldsymbol{\alpha}_1 + \cdots + \mathbb{Z}[\Lambda]\boldsymbol{\alpha}_J), \tag{7.2}$$

where the $\boldsymbol{\alpha}_j$, $1 \le j \le J$, are as above.

As a consequence of the containment (7.2), one has the following *rigid structure property* of $\mathcal{T}$ (we recall that $\Xi(\mathcal{T})$ stands for the set of return vectors of the tiling $\mathcal{T}$ as defined in (5.2)):

$$\Xi(\mathcal{T}) \subset \rho(\mathbb{Z}[\Lambda]\boldsymbol{\alpha}_1 + \cdots + \mathbb{Z}[\Lambda]\boldsymbol{\alpha}_J). \tag{7.3}$$

Theorem 7.3 ([78]). *Let $\mathcal{T}$ be a primitive FLC substitution tiling of $\mathbb{R}^n$ with expansion Λ satisfying* [A1] *and* [A2]. *Then the following are equivalent.*

(i) $\mathrm{Spec}(\Lambda)$ *is a Pisot family: if $\lambda \in \mathrm{Spec}(\Lambda)$ is of multiplicity k and λ' is a conjugate of λ with $|\lambda'| \ge 1$, then $\lambda' \in \mathrm{Spec}(\Lambda)$ and λ' has multiplicity at least k.*
(ii) *The set of eigenvalues of $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ is relatively dense in $\mathbb{R}^n$.*
(iii) *$(\Omega_{\mathcal{T}}, \mathbb{R}^n)$ is not weakly mixing* (i.e., *it has a non-zero eigenvalue*).
(iv) *$\Xi(\mathcal{T})$ is a Meyer set.*

Under the assumptions of Theorem 7.3, it is shown in [28] that the coincidence rank of the underlying substitution is finite, establishing sufficiency in Conjecture 7.1 in case [A1] and [A2] hold.

Proof. Let us sketch the proof of Theorem 7.3. The proof of $(i) \Rightarrow (ii)$ is based on the rigid structure property of $\mathcal{T}$; we sketch the argument in a simple case. Assume

that Λ has only real eigenvalues, λ_i, of multiplicity one and $\mathcal{T}$ has no translational periods. There is then a vector $\boldsymbol{\alpha} \in \mathbb{R}^n$ such that $\mathcal{C} \subset \mathbb{Z}[\Lambda]\boldsymbol{\alpha}$. Observe that

$$\Xi(\mathcal{T}) \subset \mathcal{C} - \mathcal{C} \subset \mathbb{Z}[\Lambda]\boldsymbol{\alpha}.$$

The set of control points $\mathcal{C}$ is relatively dense, consequently the vector $\boldsymbol{\alpha} := [a_1, \ldots, a_n]^T$ must have all non-zero coordinates. Consider now the vector $\boldsymbol{\beta} := [a_1^{-1}, \ldots, a_n^{-1}]^T$. We claim that the set $\{\Lambda^j\boldsymbol{\beta}\}_{j=0}^{n-1}$ is contained in the set of eigenvalues of $(\Omega_{\mathcal{T}}, \mathbb{R}^n)$. This set is a basis for $\mathbb{R}^n$ (over $\mathbb{R}$), and, since the set of eigenvalues forms an additive group, we may conclude that the eigenvalues are relatively dense. For $\mathbf{x} = \Lambda^i\boldsymbol{\alpha},\ \boldsymbol{\gamma} = \Lambda^j\boldsymbol{\beta}$ we have:

$$\langle \Lambda^l \mathbf{x}, \boldsymbol{\gamma} \rangle = \langle \Lambda^{l+i}\boldsymbol{\alpha}, \Lambda^j\boldsymbol{\beta} \rangle = \sum_{k=1}^{n} \lambda_k^{l+i+j} \xrightarrow[l\to\infty]{} 0 \pmod{\mathbb{Z}}. \tag{7.4}$$

The convergence follows from the Pisot family property: the numbers $\lambda_1, \ldots, \lambda_n$ are all roots of the same integer polynomial, and the 'missing roots' in (7.4) are all less than one in modulus. The sum of $(l+i+j)$th powers over all roots of an integer polynomial is an integer, yielding (7.4). It follows from (7.4) that

$$\lim_{l\to\infty} e^{2\pi i \langle \Lambda^l \mathbf{x}, \gamma \rangle} = 1 \quad \text{for all } \mathbf{x} \in \Xi(\mathcal{T}), \tag{7.5}$$

which means that $\boldsymbol{\gamma} \in \mathbb{R}^d$ is an eigenvalue from [113].

The proof of (ii) $\Rightarrow$ (iii) is trivial. The necessity (iii) $\Rightarrow$ (i) was proved by Robinson [99] in a more general case; it is a consequence of [113]. For the equivalence of (iv) with the rest, see [77]. □

7.3. Examples

The tile substitutions of the following two examples are primitive and two-dimensional but without FLC. Under the additional assumptions [A1] and [A2] on a substitution, FLC implies the rigid structure property. Example 7.5 shows that the converse is not true, hence substitution tilings with the rigid structure property constitute a strictly larger class than substitution tilings with FLC. The rigid structure property (7.3) is easy to check from the digit sets of the substitution and one can study various spectral properties of tiling systems in this larger class.

Example 7.4 ([68, 78]). Consider the substitution tiling $\mathcal{T}$ in $\mathbb{R}^2$ with a single prototile T and expansion $\Lambda = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}$ such that

$$\Lambda T = \bigcup_{d\in\mathcal{D}} (T + d)$$

where

$$\mathcal{D} = \{(0,-1), (0,0), (0,1), (-1,-1), (-1,0), (-1,1), (1,-1+a), (1,a), (1,1+a)\}$$

is a digit set and $a \in \mathbb{R}$ is irrational. Note that

$$\Xi(\mathcal{T}) \subset \mathbb{Z}[\Lambda](1,0) + \mathbb{Z}[\Lambda](0,1) + \mathbb{Z}[\Lambda](0,a)$$

FIGURE 2. The non-FLC Frank–Robinson substitution.

and $\mathbb{Z}[\Lambda](1,0)+\mathbb{Z}[\Lambda](0,1)+\mathbb{Z}[\Lambda](0,a)$ is the minimal module over $\mathbb{Z}[\Lambda]$ containing $\Xi(\mathcal{T})$. Thus $\mathcal{T}$ does not have the rigid structure property.

Example 7.5 ([53]). The Frank–Robinson substitution is pictured in Figure 2, in which b, the scalar expansion, is the largest (and non-Pisot) root of $x^2-x-3=0$. The digit sets $\mathcal{D}_{ij}$ satisfy $\mathcal{D}_{ij}\subset\mathbb{Z}[b](1,0)+\mathbb{Z}[b](0,1)$. Hence

$$\Xi(\mathcal{T})\subset\mathbb{Z}[b](1,0)+\mathbb{Z}[b](0,1)$$

and the rigid structure property holds. See also [16, Example 5.8] for a discussion of the Frank–Robinson tiling and see [94] for more examples of non-FLC 'fusion' tilings.

7.4. The Pisot property and hyperbolicity

In one dimension, the construction of tile substitutions is purely combinatorial. In higher dimensions, there is geometry to deal with (tile shapes) making it much more difficult to construct examples. And then it is more time-consuming to check their spectral properties. For these reasons, possible extensions of the Pisot Substitution Conjecture to higher-dimensional tile substitutions have not been well vetted. We nevertheless discuss the possibilities. The most straightforward route is:

Conjecture 7.6. Suppose that Φ is an n-dimensional primitive FLC tile substitution whose expansion Λ has the Pisot property. If the characteristic polynomial of the substitution matrix M_Φ is irreducible, then $(\Omega_\Phi,\mathbb{R}^n)$ has pure discrete spectrum.

It seems to be a fairly stringent requirement, in higher dimensions, that the characteristic polynomial of the substitution matrix M_Φ be irreducible and the irreducibility hypothesis suffers the same unnaturality as in one dimension (a different 'presentation' of the tiling space may well change the irreducibility of the characteristic polynomial of the substitution matrix). By considering a homological condition instead, we will at least be led to an interesting connection with hyperbolic dynamics. Suppose that Φ is an n-dimensional primitive FLC tile substitution with expansion Λ. Let us say that Φ is *unimodular* if every eigenvalue of Λ is an algebraic unit, and *hyperbolic* if no eigenvalue of Λ has an algebraic conjugate on the unit circle. Let $\langle\Xi(\Phi)\rangle$, called the *module of generalized return vectors* of Φ, be the additive subgroup of $\mathbb{R}^n$ generated by the return vectors $\Xi(\Phi)$ (that is, $\Xi(\Phi):=\Xi(\mathcal{T})$ for any $\mathcal{T}\in\Omega_\Phi$). If $v\in\Xi(\Phi)$, say $T,T-v\in\mathcal{T}\in\Omega_\Phi$, then $\Phi(T),\Phi(T)-\Lambda v\subset\Phi(\mathcal{T})\in\Omega_\Phi$, so that $\Lambda v\in\Xi(\Phi)$. Thus Λ induces a homomorphism $\Lambda:\langle\Xi(\Phi)\rangle\to\langle\Xi(\Phi)\rangle$. If Φ is unimodular, then $\langle\Xi(\Phi)\rangle$ is a finitely generated free $\mathbb{Z}$-module: let $D=D(\Phi):=\ \mathrm{rank}(\langle\Xi(\Phi)\rangle)$. The following is proved

in [27] using the global shadowing technique in hyperbolic dynamics pioneered by Franks [54].

Theorem 7.7 ([27]). *If Φ is a primitive n-dimensional FLC unimodular hyperbolic tile substitution with linear expansion Λ, and module of generalized return vectors $\langle \Xi(\Phi) \rangle$, there is a continuous and boundedly finite-to-one map $G : \Omega_\Phi \to \mathbb{T}^D$ so that $G \circ \Phi = F_A \circ G$. Here A is an integral unimodular hyperbolic matrix representing $\Lambda : \langle \Xi(\Phi) \rangle \to \langle \Xi(\Phi) \rangle$, and $F_A : \mathbb{T}^D \to \mathbb{T}^D$ is the hyperbolic toral automorphism associated with A. The map G is topologically essential in that the homomorphism $G^* \colon H^1(\mathbb{T}^D) \to H^1(\Omega_\phi)$ induced on first cohomology is injective and there is $r \in \mathbb{N}$ so that G is a.e. r-to-1 with respect to the translation invariant measure μ on Ω_Φ.*

If the unstable dimension of A (the sum of the dimensions of all the generalized eigenspaces of A corresponding to eigenvalues of modulus greater than one) is greater than n, then G does not semi-conjugate the translation action on Ω_Φ with a Kronecker action by $\mathbb{R}^n$ on $\mathbb{T}^D$: the map $v \mapsto G(\mathcal{T} - v)$ wiggles around in a very jagged (probably nowhere smooth) manner in the unstable manifold of $G(\mathcal{T})$ in $\mathbb{T}^D$. If the unstable dimension of A equals n, leaving no room for such wiggling, it seems plausible that G would also semi-conjugate $\mathbb{R}^n$-actions. This appears to be what happens if Λ has the Pisot property (note that if Λ has the Pisot property, then Φ is hyperbolic).

Let us partition the eigenvalues of Λ into conjugacy classes $\mathcal{F}_i$ and, for each i, let d_i be the algebraic degree of the elements of $\mathcal{F}_i$ and let J_i be the maximum multiplicity (as eigenvalues of Λ) of the elements of $\mathcal{F}_i$. The *generalized degree* of Λ is

$$d(\Lambda) := \sum J_i d_i.$$

One can show that $D(\Phi) \geq d(\Lambda)$. Thus, if Λ has the Pisot property, $D(\Phi) = d(\Lambda)$ forces the unstable dimension of A to be n, the dimension of the substitution. We know of no Pisot property substitution with $D(\Phi) > d(\Lambda)$.

Theorem 7.8 ([27]). *Suppose that Φ is a primitive FLC n-dimensional unimodular substitution whose expansion has the Pisot property. If the rank D of the module $\langle \Xi(\Phi) \rangle$ of generalized return vectors for Φ equals the generalized degree $d(\Lambda)$ of Λ, then the map $G : \Omega_\Phi \to \mathbb{T}^D$ of Theorem 7.7 is surjective and also semi-conjugates the translation action on Ω_Φ with a Kronecker action by $\mathbb{R}^n$ on $\mathbb{T}^D$. Furthermore, G is the maximal equicontinuous factor map for $(\Omega_\Phi, \mathbb{R}^n)$. That is, G=g, and hence $r = cr(\Phi) < \infty$.*

In the proof of Theorem 7.7 a *universal abelian cover*, $\tilde{\Omega}_\Phi$, is constructed and the map Φ on Ω_Φ is lifted to $\tilde{\Phi}$ on $\tilde{\Omega}_\Phi$. It is shown that the structure relation for G, denoted $\sim_{gs}$, is the global shadowing relation: $G(\mathcal{T}) = G(\mathcal{T}')$ if and only if $\mathcal{T} \sim_{gs} \mathcal{T}'$ if and only if there are $\tilde{\mathcal{T}}, \tilde{\mathcal{T}}' \in \tilde{\Omega}_\Phi$, lying over $\mathcal{T}, \mathcal{T}'$ so that the distance between $\tilde{\Phi}^k(\tilde{\mathcal{T}})$ and $\tilde{\Phi}^k(\tilde{\mathcal{T}}')$ is uniformly bounded for $k \in \mathbb{Z}$. It is a fundamental theorem of Veech [119] that the structure relation for the maximal equicontinuous

factor map g is regional proximality: tilings $\mathcal{T}, \mathcal{T}' \in \Omega_\Phi$ are *regionally proximal*, $\mathcal{T} \sim_{rp} \mathcal{T}'$, if and only if, for each $\epsilon > 0$ there are $\mathcal{S}, \mathcal{S}' \in \Omega_\Phi$ and $v \in \mathbb{R}^n$ so that: (i) $d(\mathcal{T}, \mathcal{S}) < \epsilon$, (ii) $d(\mathcal{T}', \mathcal{S}') < \epsilon$, and (iii) $d(\mathcal{S} - v, \mathcal{S}' - v) < \epsilon$. (For a general discussion of regional proximality in tiling spaces see [28].) In the context of Theorem 7.7 one can show that $\mathcal{T} \sim_{gs} \mathcal{T}' \implies \mathcal{T} \sim_{rp} \mathcal{T}'$. Under the hypotheses of Theorem 7.8, the global shadowing and regional proximal relations are the same.

The rank of $\langle \Xi(\Phi) \rangle$ is bounded above by the dimension of $H^1(\Omega_\Phi)$ so, in light of Theorem 7.8, a natural generalization to higher dimensions of the one-dimensional Homological Pisot Substitution Conjecture is:

Conjecture 7.9. Suppose that Φ is an n-dimensional primitive FLC unimodular tile substitution whose expansion Λ has the Pisot property and has generalized degree $d(\Lambda)$. If $\dim(H^1(\Omega_\Phi)) = d(\Lambda)$, then $(\Omega_\Phi, \mathbb{R}^n)$ has pure discrete spectrum.

An extension of this to the more general hyperbolic setting is:

Conjecture 7.10. Suppose that Φ is a primitive FLC unimodular and hyperbolic tile substitution and that the induced isomorphism $\Phi^*\colon H^1(\Omega_\phi) \to H^1(\Omega_\phi)$ is hyperbolic on $H^1(\Omega_\Phi)$. There is then a continuous, μ-a.e. one-to-one, topologically essential map $G : \Omega_\Phi \to \mathbb{T}^D$, $D = \dim(H^1(\Omega_\Phi))$, and a hyperbolic toral automorphism $F_A : \mathbb{T}^D \to \mathbb{T}^D$, with $G \circ \Phi = F_A \circ G$.

Just as it is rare that characteristic polynomials of substitution matrices are irreducible in higher dimensions, it is unusual for the first cohomology of the tiling space to have the minimal dimension required to accommodate the generalized degree of the expansion. This limits the range of the above conjectures. They can be strengthened, as can Conjectures 4.4 and 4.5, by replacing the cohomology of Ω_Φ by its *essential cohomology*, which does not see, for example, the contributions of asymptotic cycles (see [27]). The resulting conjectures imply the non-homological conjectures in the unimodular case.

Acknowledgment

We would like to express our thanks to the anonymous referee for numerous and valuable suggestions. We also would like to thank B. Loridant for his careful reading.

References

[1] S. Akiyama, "Pisot numbers and greedy algorithm," in *Number theory (Eger, 1996)*, pp. 9–21, Berlin: de Gruyter, 1998.

[2] S. Akiyama, "Cubic Pisot units with finite beta expansions," in *Algebraic number theory and Diophantine analysis* (*Graz,* 1998), pp. 11–26, Berlin: de Gruyter, 2000.

[3] S. Akiyama, "On the boundary of self affine tilings generated by Pisot numbers," *J. Math. Soc. Japan*, vol. 54, no. 2, pp. 283–308, 2002.

[4] S. Akiyama, "Positive finiteness of number systems," in *Number theory*, vol. 15 of *Dev. Math.*, pp. 1–10, New York: Springer, 2006.

[5] S. Akiyama, "Pisot number system and its dual tiling," in *Physics and theoretical computer science*, vol. 7 of *NATO Secur. Sci. Ser. D Inf. Commun. Secur.*, pp. 133–154, Amsterdam: IOS, 2007.

[6] S. Akiyama and J.-Y. Lee, "Algorithm for determining pure pointedness of self-affine tilings," *Adv. Math.*, vol. 226, no. 4, pp. 2855–2883, 2011.

[7] S. Akiyama and J.-Y. Lee, "Determining pure discrete spectrum for some self-affine tilings," *Discrete Math. Theor. Comput. Sci.*, vol. 16, no. 3, pp. 305–316, 2014.

[8] S. Akiyama and K. Scheicher, *From number systems to shift radix systems*, Nihonkai Math. J. **16** (2005), no. 2, 95–106.

[9] S. Akiyama, H. Rao, and W. Steiner, "A certain finiteness property of Pisot number systems," *J. Number Theory*, vol. 107, no. 1, pp. 135–160, 2004.

[10] S. Akiyama, G. Barat, V. Berthé, and A. Siegel, "Boundary of central tiles associated with Pisot beta-numeration and purely periodic expansions," *Monatsh. Math.*, vol. 155, no. 3-4, pp. 377–419, 2008.

[11] S. Akiyama, H. Brunotte, A. Pethő, and J.M. Thuswaldner, "Generalized radix representations and dynamical systems. II," *Acta Arith.*, vol. 121, no. 1, pp. 21–61, 2006.

[12] S. Akiyama, T. Borbély, H. Brunotte, A. Pethő, and J.M. Thuswaldner, "Generalized radix representations and dynamical systems. I," *Acta Math. Hungar.*, vol. 108, no. 3, pp. 207–238, 2005.

[13] J.E. Anderson and I.F. Putnam, "Topological invariants for substitution tilings and their associated C^*-algebras," *Ergodic Theory Dynam. Systems*, vol. 18, no. 3, pp. 509–537, 1998.

[14] P. Arnoux and S. Itô, "Pisot substitutions and Rauzy fractals," *Bull. Belg. Math. Soc. Simon Stevin*, vol. 8, no. 2, pp. 181–207, 2001. Journées Montoises (Marne-la-Vallée, 2000).

[15] P. Arnoux, M. Furukado, E. Harriss, and S. Itô, "Algebraic numbers, free group automorphisms and substitutions on the plane," *Trans. Amer. Math. Soc.*, vol. 363, no. 9, pp. 4651–4699, 2011.

[16] M. Baake and U. Grimm, *Aperiodic order. Vol.* 1, vol. 149 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2013.

[17] M. Baake and D. Lenz, "Dynamical systems on translation bounded measures: pure point dynamical and diffraction spectra," *Ergodic Theory Dynam. Systems*, vol. 24, no. 6, pp. 1867–1893, 2004.

[18] M. Baake and D. Lenz, "Deformation of Delone dynamical systems and pure point diffraction," *J. Fourier Anal. Appl.*, vol. 11, no. 2, pp. 125–150, 2005.

[19] M. Baake and R.V. Moody, "Weighted Dirac combs with pure point diffraction," *J. Reine Angew. Math.*, vol. 573, pp. 61–94, 2004.

[20] M. Baake, D. Lenz, and R. Moody, "Characterization of model sets by dynamica. systems," *Ergodic Theory Dynam. Systems*, vol. 27, pp. 341–382, 2007.

[21] V. Baker, M. Barge, and J. Kwapisz, "Geometric realization and coincidence for reducible non-unimodular Pisot tiling spaces with an application to beta-shifts," *Ann. Inst. Fourier* (*Grenoble*), vol. 56, no. 7, pp. 2213–2248, 2006.

[22] M. Barge, "Factors of Pisot tiling spaces and the coincidence rank conjecture." *Bull. Soc. Math. France*, to appear.

[23] M. Barge, "Pure discrete spectrum for a class of one-dimensional substitution tiling systems." Preprint, 2014.

[24] M. Barge and B. Diamond, "A complete invariant for the topology of one-dimensional substitution tiling spaces," *Ergod. Th. and Dynam. Sys.*, vol. 21, pp. 1333–1358, 2001.

[25] M. Barge and B. Diamond, "Proximality in Pisot tiling spaces," *Fund. Math.*, vol. 194, pp. 191–238, 2007.

[26] M. Barge and B. Diamond, "Cohomology in one-dimensional substitution tiling spaces," *Proc. Amer. Math. Soc.*, vol. 136, no. 6, pp. 2183–2191, 2008.

[27] M. Barge and J.-M. Gambaudo, "Geometric realization of substitution tiling spaces," *Ergodic Theory Dynam. Systems*, vol. 34, no. 2, pp. 443–468, 2014.

[28] M. Barge and J. Kellendonk, "Proximality and pure point spectrum for tiling dynamical systems," *Michigan Math. J.*, vol. 62, no. 4, pp. 793–822, 2013.

[29] M. Barge and J. Kwapisz, "Elements of the theory of unimodular Pisot substitutions with an application to β-shifts," in *Algebraic and topological dynamics*, vol. 385 of *Contemporary Mathematics*, pp. 89–99, Amer. Math. Soc., 2005.

[30] M. Barge and J. Kwapisz, "Geometric theory of unimodular Pisot substitutions," *Amer. J. Math.*, vol. 128, no. 5, pp. 1219–1282, 2006.

[31] M. Barge and R. Swanson, "Rigidity in one-dimensional tiling spaces," *Topology Appl.*, vol. 154, no. 17, pp. 3095–3099, 2007.

[32] M. Barge, S. Štimac, and R.F. Williams, "Pure discrete spectrum in substitution tiling spaces," *Discrete Contin. Dyn. Syst.*, vol. 33, no. 2, pp. 579–597, 2013.

[33] M. Barge, H. Bruin, L. Jones, and L. Sadun, "Homological Pisot substitutions and exact regularity," *Israel J. Math.*, vol. 188, pp. 281–300, 2012.

[34] V. Berthé and M. Rigo, *Combinatorics, Automata and Number Theory*, vol. 135 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2010.

[35] V. Berthé and A. Siegel, "Tilings associated with beta-numeration and substitutions," *Integers*, vol. 5, 2005.

[36] V. Berthé and A. Siegel, "Purely periodic β-expansions in the Pisot non-unit case," *J. Number Theory*, vol. 127, no. 2, pp. 153–172, 2007.

[37] V. Berthé, D. Frettlöh, and V. Sirvent, "Selfdual substitutions in dimension one," *European J. Combin.*, vol. 33, no. 6, pp. 981–1000, 2012.

[38] V. Berthé, T. Jolivet, and A. Siegel, "Substitutive Arnoux–Rauzy sequences have pure discrete spectrum," *Unif. Distrib. Theory*, vol. 7, no. 1, pp. 173–197, 2012.

[39] V. Berthé, A. Siegel, and J.M. Thuswaldner, "Substitutions, Rauzy fractals, and tilings," in *Combinatorics, Automata and Number Theory*, vol. 135 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, 2010.

[40] V. Berthé, J. Bourdon, T. Jolivet, and A. Siegel, "Generating discrete planes with substitutions," in *WORDS* (J. Karhumäki, A. Lepistö, and L.Q. Zamboni, eds.), vol. 8079 of *Lecture Notes in Computer Science*, pp. 58–70, Springer, 2013.

[41] V. Berthé, J. Bourdon, T. Jolivet, and A. Siegel, "A combinatorial approach to products of Pisot substitutions." *Ergodic Theory Dynam. Systems*, to appear.

[42] E. Bombieri and J.E. Taylor, "Which distributions of matter diffract? an initial investigation," in *International Workshop on Aperiodic Crystals* (D. Gratias and L. Michel, eds.), vol. 47, pp. 19–28, 1986.

[43] E. Bombieri and J.E. Taylor, "Quasicrystals, tilings, and algebraic number theory: some preliminary connections," in *The legacy of Sonya Kovalevskaya* (*Cambridge, Mass., and Amherst, Mass.,* 1985), vol. 64 of *Contemp. Math.*, pp. 241–264, Providence, RI: Amer. Math. Soc., 1987.

[44] R. Bowen, "Markov partitions for Axiom-A diffeomorphisms," *Proc. Amer. Math. Soc.*, vol. 92, pp. 725–, 1970.

[45] A. Clark and L. Sadun, "When size matters: subshifts and their related tiling spaces," *Ergodic Theory Dynam. Systems*, vol. 23, pp. 1043–1057, 2003.

[46] A. Clark and L. Sadun, "When shape matters: deformations of tiling spaces," *Ergodic Theory Dynam. Systems*, vol. 26, pp. 69–86, 2006.

[47] F.M. Dekking, "The spectrum of dynamical systems arising from substitutions of constant length," *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, vol. 41, no. 3, pp. 221–239, 1977/78.

[48] S. Dworkin, "Spectral theory and x-ray diffraction," *J. Math. Phys.*, vol. 34, no. 7, pp. 2965–2967, 1993.

[49] H. Ei and S. Itô, "Tilings from some non-irreducible, Pisot substitutions," *Discrete Math. and Comp. Science*, vol. 8, no. 1, pp. 81–122, 2005.

[50] H. Ei, S. Itô, and H. Rao, "Atomic surfaces, tilings and coincidences. II. Reducible case," *Université de Grenoble. Annales de l'Institute Fourier*, vol. 56, no. 7, pp. 2285–2313, 2006.

[51] M. Einsiedler and T. Ward, *Ergodic theory with a view towards number theory*, vol. 259 of *Graduate Texts in Mathematics*. Springer-Verlag London, Ltd., London, 2011.

[52] D.-J. Feng, M. Furukado, S. Itô, and J. Wu, "Pisot substitutions and the Hausdorff dimension of boundaries of atomic surfaces," *Tsukuba J. Math.*, vol. 30, no. 1, pp. 195–223, 2006.

[53] N.P. Frank and E.A. Robinson, Jr., "Generalized β-expansions, substitution tilings, and local finiteness," *Trans. Amer. Math. Soc.*, vol. 360, no. 3, pp. 1163–1177 (electronic), 2008.

[54] J. Franks, "Anosov diffeomorphisms on tori," *Trans. Amer. Math. Soc.*, vol. 145, pp. 117–124, 1969.

[55] D. Frettlöh, "Self-dual tilings with respect to star-duality," *Theoret. Comput. Sci.*, vol. 391, no. 1-2, pp. 39–50, 2008.

[56] D. Frettlöh and B. Sing, "Computing modular coincidences for substitution tilings and point sets," *Discrete Comput. Geom.*, vol. 37, no. 3, pp. 381–407, 2007.

[57] C. Frougny and B. Solomyak, "Finite beta-expansions," *Ergodic Theory Dynam. Systems*, vol. 12, no. 4, pp. 713–723, 1992.

[58] J.-P. Gazeau and J.-L. Verger-Gaugry, "Geometric study of the beta-integers for a Perron number and mathematical quasicrystals," *J. Théor. Nombres Bordeaux*, vol. 16, no. 1, pp. 125–149, 2004.

[59] W.J. Gilbert, "Radix representations of quadratic fields," *J. Math. Anal. Appl.*, vol. 83, no. 1, pp. 264–274, 1981.

[60] J.-B. Gouéré, "Quasicrystals and almost periodicity," *Comm. Math. Phys.*, vol. 255, no. 3, pp. 655–681, 2005.

[61] A. Hof, "On diffraction by aperiodic structures," *Comm. Math. Phys.*, vol. 169, no. 1, pp. 25–43, 1995.

[62] A. Hof, "Diffraction by aperiodic structures," in *The mathematics of long-range aperiodic order* (*Waterloo, ON,* 1995), vol. 489 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pp. 239–268, Dordrecht: Kluwer Acad. Publ., 1997.

[63] M.I. Hollander, *Linear numeration systems, finite beta expansions, and discrete spectrum of substitution dynamical systems*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.) – University of Washington.

[64] S. Itô and M. Ohtsuki, "Parallelogram tilings and Jacobi-Perron algorithm," *Tokyo J. Math.*, vol. 17, no. 1, pp. 33–58, 1994.

[65] S. Itô and H. Rao, "Atomic surfaces, tilings and coincidences I. Irreducible case," *Israel J. Math.*, vol. 153, pp. 129–156, 2006.

[66] S. Itô and Y. Takahashi, "Markov subshifts and realization of β-expansions," *J. Math. Soc. Japan*, vol. 26, pp. 33–55, 1974.

[67] T. Kamae, "A topological invariant of substitution minimal sets," *J. Math. Soc. Japan*, vol. 24, pp. 285–306, 1972.

[68] R. Kenyon, "Self-replicating tilings," in *Symbolic dynamics and its applications* (*New Haven, CT,* 1991), vol. 135 of *Contemp. Math.*, pp. 239–263, Providence, RI: Amer. Math. Soc., 1992.

[69] R. Kenyon and B. Solomyak, "On the characterization of expansion maps for self-affine tilings," *Discrete Comput. Geom.*, vol. 43, no. 3, pp. 577–593, 2010.

[70] R. Kenyon and A. Vershik, "Arithmetic construction of sofic partitions of hyperbolic toral automorphisms," *Ergodic Theory Dynam. Systems*, vol. 18, no. 2, pp. 357–372, 1998.

[71] J. Kwapisz, "Rigidity and mapping class group for abstract tiling spaces," *Ergodic Theory Dynam. Systems*, vol. 31, no. 6, pp. 1745–1783, 2011.

[72] J. Kwapisz, "Expansions of self-affine tilings are integral Perron (the non-diagonalizable case)." Preprint, 2013.

[73] J.C. Lagarias, "Meyer's concept of quasicrystal and quasiregular sets," *Comm. Math. Phys.*, vol. 179, no. 2, pp. 365–376, 1996.

[74] J.C. Lagarias and Y. Wang, "Substitution Delone sets," *Discrete Comput. Geom.*, vol. 29, no. 2, pp. 175–209, 2003.

[75] J.-Y. Lee, "Substitution Delone sets with pure point spectrum are inter-model sets," *J. Geom. Phys.*, vol. 57, no. 11, pp. 2263–2285, 2007.

[76] J.-Y. Lee and R.V. Moody, "Lattice substitution systems and model sets," *Discrete Comput. Geom.*, vol. 25, no. 2, pp. 173–201, 2001.

[77] J.-Y. Lee and B. Solomyak, "Pure point diffractive substitution Delone sets have the Meyer property," *Discrete Comput. Geom.*, vol. 39, no. 1-3, pp. 319–338, 2008.

[78] J.-Y. Lee and B. Solomyak, "Pisot family self-affine tilings, discrete spectrum, and the Meyer property," *Discrete Contin. Dyn. Syst.*, vol. 32, no. 3, pp. 935–959, 2012.

[79] J.-Y. Lee, D. Lenz, and B. Sing, "Strongly almost periodic measures and diffraction theory for modulated structures." Preprint, 2015.

[80] J.-Y. Lee, R.V. Moody, and B. Solomyak, "Pure point dynamical and diffraction spectra," *Ann. Henri Poincaré*, vol. 3, no. 5, pp. 1003–1018, 2002.

[81] J.-Y. Lee, R.V. Moody, and B. Solomyak, "Consequences of pure point diffraction spectra for multiset substitution systems," *Discrete Comput. Geom.*, vol. 29, no. 4, pp. 525–560, 2003.

[82] D. Lenz and N. Strungaru, "Pure point spectrum for measure dynamical systems on locally compact abelian groups," *J. Math. Pures Appl.* (9), vol. 92, no. 4, pp. 323–341, 2009.

[83] A.N. Livshits, "On the spectra of adic transformations of Markov compact sets," *Uspekhi Mat. Nauk*, vol. 42, no. 3(255), pp. 189–190, 1987. English translation: *Russian Math. Surveys* 42(3): 222–223, 1987.

[84] B. Loridant, A. Messaoudi, P. Surer, and J.M. Thuswaldner, "Tilings induced by a class of cubic Rauzy fractals," *Theoret. Comput. Sci.*, vol. 477, pp. 6–31, 2013.

[85] M. Lothaire, *Algebraic combinatorics on words*, vol. 90 of *Encyclopedia of Mathematics and its Applications*. Cambridge: Cambridge University Press, 2002.

[86] B.F. Martensen, "Generalized balanced pair algorithm," *Topology Proc.*, vol. 28, no. 1, pp. 163–178, 2004. Spring Topology and Dynamical Systems Conference.

[87] P. Michel, "Coincidence values and spectra of substitutions," *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, vol. 42, no. 3, pp. 205–227, 1978.

[88] R.V. Moody, "Meyer sets and their duals," in *The mathematics of long-range aperiodic order* (*Waterloo, ON,* 1995), vol. 489 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pp. 403–441, Dordrecht: Kluwer Acad. Publ., 1997.

[89] R.V. Moody, "Model sets: a survey," in *From Quasicrystals to More Complex Systems*, Springer Verlag, 2000. F. Axel, F. Dénoyer, and J.-P. Gazeau, Centre de physique Les Houches.

[90] H.M. Morse, "Recurrent geodesics on a surface of negative curvature," *Trans. Amer. Math. Soc.*, vol. 22, no. 1, pp. 84–100, 1921.

[91] H.M. Morse and G.A. Hedlund, "Symbolic Dynamics," *Amer. J. Math.*, vol. 60, no. 4, pp. 815–866, 1938.

[92] C. Pisot, "La répartition modulo 1 et nombres algébriques," *Ann. Sc. Norm. Super. Pisa*, vol. II, pp. 205–248, 1938.

[93] B. Praggastis, "Numeration systems and Markov partitions from self-similar tilings," *Trans. Amer. Math. Soc.*, vol. 351, no. 8, pp. 3315–3349, 1999.

[94] N. Priebe Frank and L. Sadun, "Fusion tilings with infinite local complexity," *Topology Proc.*, vol. 43, pp. 235–276, 2014.

[95] N. Pytheas Fogg, *Substitutions in Dynamics, Arithmetics and Combinatorics*, vol. 1794 of *Lecture Notes in Mathematics*. Springer-Verlag, 2002. Ed. by V. Berthé and S. Ferenczi and C. Mauduit and A. Siegel.

[96] M. Queffélec, *Substitution dynamical systems – spectral analysis*. Berlin: Springer-Verlag, 1987. Lecture Notes in Mathematics, Vol. 1294.

[97] G. Rauzy, "Nombres algébriques et substitutions," *Bull. Soc. Math. France*, vol. 110, pp. 147–178, 1982.

[98] A. Rényi, "Representations for real numbers and their ergodic properties," *Acta Math. Acad. Sci. Hungar*, vol. 8, pp. 477–493, 1957.

[99] E.A. Robinson, Jr., "Symbolic dynamics and tilings of $\mathbb{R}^d$," in *Symbolic dynamics and its applications*, vol. 60 of *Proc. Sympos. Appl. Math., Amer. Math. Soc. Providence, RI*, pp. 81–119, 2004.

[100] L. Sadun, *Topology of tiling spaces*, vol. 46 of *University Lecture Series*. Providence, RI: American Mathematical Society, 2008.

[101] M. Schlottmann, "Cut-and-project sets in locally compact abelian groups," in *Quasicrystals and discrete geometry* (*Toronto, ON,* 1995), vol. 10 of *Fields Inst. Monogr.*, pp. 247–264, Providence, RI: Amer. Math. Soc., 1998.

[102] M. Schlottmann, "Generalized model sets and dynamical systems," in *Directions in mathematical quasicrystals*, vol. 13 of *CRM Monogr. Ser.*, pp. 143–159, Amer. Math. Soc., Providence, RI, 2000.

[103] K. Schmidt, "Algebraic coding of expansive group automorphisms and two-sided beta-shifts," *Monatsh. Math.*, vol. 129, no. 1, pp. 37–61, 2000.

[104] K. Schmidt, "Quotients of $l^\infty(\mathbb{Z}, \mathbb{Z})$ and symbolic covers of toral automorphisms," in *Representation theory, dynamical systems, and asymptotic combinatorics*, vol. 217 of *Amer. Math. Soc. Transl. Ser. 2*, pp. 223–246, Providence, RI: Amer. Math. Soc., 2006.

[105] N. Sidorov, "Bijective and general arithmetic codings for Pisot toral automorphisms," *J. Dynamics and Control Systems*, vol. 7, pp. 447–472, 2001.

[106] N. Sidorov, "Arithmetic dynamics," in *Topics in dynamics and ergodic theory*, vol. 310 of *London Math. Soc. Lecture Note Ser.*, pp. 145–189, Cambridge: Cambridge Univ. Press, 2003.

[107] A. Siegel, "Représentation des systèmes dynamiques substitutifs non unimodulaires," *Ergodic Theory Dynam. Systems*, vol. 23, no. 4, pp. 1247–1273, 2003.

[108] A. Siegel and J.M. Thuswaldner, "Topological properties of Rauzy fractals," *Mém. Soc. Math. Fr.* (*N.S.*), no. 118, p. 140, 2009.

[109] B. Sing, *Pisot Substitutions and Beyond*. PhD thesis, Universität Bielefeld, 2007.

[110] B. Sing, "Modulated quasicrystals," *Z. Kristallographie*, vol. 223, no. 11-12, pp. 765–769, 2008.

[111] V.F. Sirvent and B. Solomyak, "Pure discrete spectrum for one-dimensional substitution systems of Pisot type," *Canad. Math. Bull.*, vol. 45, no. 4, pp. 697–710, 2002. Dedicated to Robert V. Moody.

[112] B. Solomyak, "A dynamical system with a discrete spectrum," *Uspekhi Mat. Nauk*, vol. 41, no. 2(248), pp. 209–210, 1986.

[113] B. Solomyak, "Dynamics of self-similar tilings," *Ergod. Th. & Dynam. Sys.*, vol. 17, pp. 695–738, 1997.

[114] B. Solomyak, "Nonperiodicity implies unique composition for self-similar translationally finite tilings," *Discrete Comput. Geom.*, vol. 20, pp. 265–279, 1998.

[115] B. Solomyak, "Eigenfunctions for substitution tiling systems," in *Probability and number theory – Kanazawa* 2005, vol. 49 of *Adv. Stud. Pure Math.*, pp. 433–454, Tokyo: Math. Soc. Japan, 2007.

[116] N. Strungaru, "Almost periodic measures and long-range order in Meyer sets," *Discrete Comput. Geom.*, vol. 33, no. 3, pp. 483–505, 2005.

[117] W. Thurston, "Groups, tilings and finite state automata." AMS Colloquium Lecture Notes, 1989.

[118] S. van Smaalen, *Incommensurate Crystallography.* Oxford: Oxford University Press, 2007.

[119] W.A. Veech, "The equicontinuous structure relation for minimal abelian transformation groups," *Amer. J. Math.*, vol. 90, pp. 723–732, 1968.

[120] A. Vince, "Digit tiling of Euclidean space," in *Directions in mathematical quasicrystals*, vol. 13 of *CRM Monogr. Ser.*, pp. 329–370, Amer. Math. Soc., Providence, RI, 2000.

[121] P. Walters, *An introduction to ergodic theory.* New York: Springer-Verlag, 1982.

S. Akiyama
Institute of Mathematics, University of Tsukuba
1-1-1 Tennodai, Tsukuba, Ibaraki
350-8571 Japan
e-mail: akiyama@math.tsukuba.ac.jp

M. Barge
Department of Mathematical Sciences
Montana State University
Bozeman, MT 59717-0240, USA
e-mail: barge@math.montana.edu

V. Berthé
Laboratoire d'Informatique Algorithmique:
Fondements et Applications,
Université Paris Diderot, Paris 7 – Case 7014
F-75205 Paris Cedex 13, France
e-mail: berthe@liafa.univ-paris-diderot.fr

J.-Y. Lee
Dept. of Math. Edu.
Catholic Kwandong University
24, 579 Bun-gil, Beomil-ro, Gangneung
Gangwon 210-701, Korea
e-mail: jylee@cku.ac.kr

A. Siegel
CNRS – Institut de Recherche en Informatique
et Systèmes Aléatoires
Campus de Beaulieu
F-35042 Rennes Cedex, France
e-mail: anne.siegel@irisa.fr

Progress in Mathematics, Vol. 309, 73–104
© 2015 Springer Basel

Cohomology of Hierarchical Tilings

Lorenzo Sadun

Abstract. We go over different versions of tiling cohomology (Čech, pattern-equivariant, PV, quotient) with emphasis on the inverse limit constructions used to compute these cohomologies. We then consider the uses of tiling cohomology to distinguish spaces, to understand deformations, and to help understand maps between tiling spaces. The emphasis of this chapter is on substitution tilings and their generalizations, but the underlying ideas apply equally well to cut-and-project tilings and to tilings defined by local matching rules.

Mathematics Subject Classification (2010). 37B50, also 52C22, 52C23, 55N05 and 55C25.

Keywords. Čech, cohomology, inverse limits, tiling, collar, deformation, pattern-equivariant, quotient, topology, substitution, force-the-border.

Algebraic invariants, such as homotopy groups, homology groups, and cohomology groups, are used to study topological spaces and maps between them. The best-known of these is homology. In a homology theory, we associate Abelian groups $C_k(X)$ of *chains* to a topological space X, and define a boundary operator $\partial_k : C_k(X) \to C_{k-1}(X)$ such that $\partial_{k-1} \circ \partial_k = 0$. Chains in the kernel of ∂_k are called closed, and chains in the image of ∂_{k+1} are called boundaries. The kth homology of X in this setup is $H_k(X) = \mathrm{Ker}(\partial_k)/\,\mathrm{Im}(\partial_{k+1})$. A continuous map $f : X \to Y$ induces *push-forward* maps $f_* : C_k(X) \to C_k(Y)$. (Strictly speaking there is one such map for each integer k, but they are all denoted f_*). This in turn induces a map (also denoted f_*) from $H_k(X)$ to $H_k(Y)$. Homotopic maps induce the same map on homology. Homology groups can then help us classify spaces, and the pushforward $f_* : H_k(X) \to H_k(Y)$ helps classify maps up to homotopy, and hence the relation between X and Y. There are many different homology theories, including simplicial, singular and cellular. For CW complexes they all yield isomorphic groups, so we often get lazy and speak of *the* homology of a space X without specifying the theory.

In a *cohomology* theory, we associate Abelian groups $C^k(X)$ of *cochains* to X and define a *coboundary* operator $\delta_k : C^k(X) \to C^{k+1}(X)$ such that $\delta_{k+1} \circ \delta_k = 0$. Given such a setup, the kth cohomology group of X is $H^k(X) = \mathrm{Ker}(\delta_k)/\,\mathrm{Im}(\delta_{k-1})$.

A continuous map $f : X \to Y$ induces *pullback* maps $f^* : C^k(Y) \to C^k(X)$ and $f^* : H^k(Y) \to H^k(X)$.

One way to get a cohomology theory is to start with a homology theory and dualize everything[1]. We can define $C^k(X)$ to be the dual space of $C_k(X)$, and δ_k to be the transpose of ∂_{k+1}. That is, if α is a k-cochain and c is a $(k+1)$-chain, then

$$(\delta_k \alpha)(c) := \alpha(\partial_{k+1} c), \tag{1}$$

since the boundary ∂c of c is a k-chain[2]. This is how simplicial, singular, and cellular cohomology are defined. However, there are also cohomology theories that are defined intrinsically rather than via homology. In de Rham cohomology, if X is a smooth manifold, then $C^k(X)$ is the set of k-forms on X, and δ_k is the exterior derivative. In Čech cohomology, the cochains are defined via open covers of X. Regardless of the setup, we call elements of Ker δ_k *co-closed* and elements of Im δ_{k-1} *co-exact*, and define

$$H^k(X) = (\text{Ker } \delta_k)/(\text{Im } \delta_{k-1}). \tag{2}$$

Since the 1990s, Čech cohomology has been used to study tiling spaces[3]. This began with work of Kellendonk [Kel1], and really took off after the seminal work of Anderson and Putnam [AP]. This chapter will address three essential questions, all of which have generated a host of papers: (1) What is tiling cohomology? (2) How do you compute it? (3) What is it good for? Most of this chapter is review material, but the content of Sections 3.1.1 and 3.1.2 is new and is joint work with John Hunton.

1. What is tiling cohomology?

Many algebraic invariants that are used to classify topological spaces do not work very well with tiling spaces. Tiling spaces (with finite local complexity) are "matchbox manifolds"; foliated spaces that locally look like the product of Euclidean space and a Cantor set. Tiling spaces have uncountably many path components. Most of the standard algebraic invariants are then useless, since they look at each path component separately, without regard to how the path components approximate one another. For instance, in singular homology, H_0 of a tiling space is a free group with uncountably many generators, while all higher homology groups vanish. The fundamental group and all higher homotopy groups also vanish.

To get around these difficulties, we need to employ less familiar cohomology theories, especially Čech cohomology, which is well adapted to tiling theory. In Subsection 1.1 we describe how to view tiling spaces as inverse limits. In Subsection 1.2 we describe Čech cohomology and explain how to view the cohomology

[1] This is where the prefix "co" for objects related to cohomology comes from.

[2] When the dimension of a chain or cochain is clear, we often omit the subscript from ∂_k or δ_k.

[3] All tilings in this chapter will be assumed to have finite local complexity, and in particular to have tiles that meet full-edge to full-edge. Cohomology can also be used to study tiling spaces of infinite local complexity, but both the calculations and the interpretations are more complicated.

of an inverse limit space. In Subsection 1.3 we go over *pattern-equivariant cohomology*. This is a theory, isomorphic to Čech cohomology, in which the cochains and cocycles can be viewed as functions on a single tiling. PV cohomology, described in subsection 1.4, is another reformulation of the Čech complex, only now the cochains are functions on Cantor sets. Finally, in subsection 1.5 we describe quotient cohomology, an analogue of relative cohomology that is very useful in computations.

1.1. Inverse limit spaces

Let $\Gamma^0, \Gamma^1, \Gamma^2, \ldots$ be a sequence of topological spaces, and for each $n > 0$ let let $\rho_n : \Gamma^n \to \Gamma^{n-1}$ be a continuous map. The *inverse limit* $\varprojlim(\Gamma^n, \rho_n)$ is a subset of the product space $\prod_n \Gamma_n$. It is the set of all sequences $(x_0, x_1, x_2, \ldots) \in \prod_n \Gamma^n$ such that for each $n > 0$, $\rho_n(x_n) = x_{n-1}$. The spaces Γ^n are called *approximants* to the inverse limit, since knowing $x_n \in \Gamma^n$ determines the first $n+1$ terms $(x_0, \ldots, x_n)$ in the sequence, and thus approximates the entire sequence in the product topology.

A simple example is the *dyadic solenoid* Sol_2. Each Γ_n is the circle $\mathbb{R}/\mathbb{Z}$, and each ρ_n is the doubling map. A point in $\mathrm{Sol}_2 = \varprojlim(S^1, \times 2)$ is a point x_0 on the unit circle, together with a choice between two possible preimages x_1, another choice between possible preimages x_2 of x_1, another choice of x_3, etc. Infinitely many discrete choices make a Cantor set, and Sol_2 is a Cantor set bundle over the circle.

There are many descriptions of tiling spaces as inverse limits, and we will present a few of the constructions in Section 2. If the tilings have finite local complexity, then the approximants are branched manifolds or branched orbifolds [AP, BBG, Sa1]. Even if the tilings do not have finite local complexity, it is usually possible to construct reasonable approximants. The approximants Γ^n parametrize the possible restrictions of a tiling to a ball of radius r_n, with $\lim_{n\to\infty} r_n = \infty$, and the maps ρ_n are obtained by restricting the tiling to a smaller region. A point in the inverse limit is a set of consistent instructions for tiling bigger and bigger balls around the origin, which is tantamount to a tiling of the entire plane.

1.2. Čech cohomology

The precise definition of the Čech cohomology $\check{H}^*(\Omega)$ of a topological space Ω involves the combinatorics of open covers of Ω, and how the combinatorics change with refinements of the open covers. The (complicated!) details can be found in an algebraic topology text [BT, Hat, Sa3] and need not concern us here. What *do* concern us are some standard properties of Čech cohomology.

Theorem 1.1. *If X is a CW complex, then the Čech cohomology $\check{H}^*(X)$ is naturally isomorphic, as a ring, to the singular cohomology $H^*(X)$, and also to the cellular cohomology. If X is a manifold, then the Čech cohomology with real coefficients is isomorphic to the de Rham cohomology $H^*_{dR}(X)$.*

Recall that if we have a sequence $G_0, G_1, \ldots$ of groups, and a collection of homomorphisms $\eta_n^* : G_n \to G_{n+1}$, then the *direct limit* $\varinjlim(G_n, \eta_n)$ is the disjoint union of the G_n's, modulo the relation that $x_n \in G_n$ is identified with $\eta_n(x_n) \in G_{n+1}$. Every element $x \in \varinjlim(G_n, \eta_n)$ is the equivalence class of an element of an approximating group G_n; *there are no additional elements "at infinity"*. For instance, $Z[1/2] := \varinjlim(\mathbb{Z}, \times 2)$ is isomorphic to the set of *dyadic rational* numbers whose denominators are powers of 2. The element $k \in G_n$ is associated with the rational number $k/2^n$, and $k \in G_n$ equals $2k \in G_{n+1}$ (as it must). The rational number 5/16 can be represented as $5 \in G_4$, $10 \in G_5$, or $20 \in G_6$, etc., but has no representative in G_0, G_1, G_2 or G_3.

Theorem 1.2. *If Ω is the inverse limit $\varprojlim(\Gamma^n, \rho_n)$ of a sequence of spaces Γ^n under a sequence of maps $\rho_n : \Gamma^n \to \Gamma^{n-1}$, then $\check{H}^*(\Omega)$ is isomorphic to the direct limit $\varinjlim(\check{H}^*(\Gamma^n), \rho_{n+1}^*)$.*

In other words, *all cohomology theories on a nice space are the same, and the Čech cohomology of an inverse limit is the direct limit of the Čech cohomologies of the approximants.*

This is how tiling cohomology is most frequently viewed in practice. Every element of $\check{H}^*(\Omega)$ can be represented by a class in $\check{H}^k(\Gamma^n)$ on some approximant Γ^n, and hence by a singular or cellular cochain on Γ^n. Instead of working with arbitrary open covers of the tiling space itself, we write everything in terms of the cells that compose the approximants.

As an example, consider the dyadic solenoid. $H^0(S^1) = H^1(S^1) = \mathbb{Z}$. Since ρ_n wraps the circle twice around itself, ρ_n^* is the identity on H^0 and multiplication by 2 on H^1. Thus $\check{H}^0(\mathrm{Sol}_2) = \varinjlim(\mathbb{Z}, \times 1) = \mathbb{Z}$ and $\check{H}^1(\mathrm{Sol}_2) = \varinjlim(\mathbb{Z}, \times 2) = \mathbb{Z}[1/2]$. If we view S^1 as consisting of one 0-cell and one 1-cell, then for each $m \geq n$, the element $2^{-n} \in \check{H}^1(\mathrm{Sol}_2)$ can be represented by a cochain on Γ^m that evaluates to 2^{m-n} on the 1-cell.

1.3. Pattern-equivariant cohomology

Tiling cohomology can also be understood in terms of the properties of a single tiling of $\mathbf{T} \in \Omega$. This approach, called *pattern-equivariant* (*PE*) *cohomology*, was developed by Kellendonk and Putnam [Kel2, KP] using differential forms, and extended to integer-valued cohomology in [Sa2].

Suppose that $f : \mathbb{R}^d \to \mathbb{R}$ is a smooth function. We say that f is *pattern-equivariant* (*or PE*) *with radius* R if the value of $f(x)$ depends only on what the tiling $\mathbf{T}$ looks like in a ball of radius R around x. That is, if $x, y \in \mathbb{R}^d$, and if $\mathbf{T} - x$ and $\mathbf{T} - y$ agree exactly on a ball of radius R around the origin, then $f(x)$ must equal $f(y)$. A function is called *strongly PE* if it is PE with some finite radius R. A function is *weakly PE* if it and all of its derivatives are uniform limits of strongly PE functions.

PE forms are defined similarly. Let $\Lambda^k_{PE}(\mathbf{T})$ denote the k-forms on $\mathbb{R}^d$ that are strongly PE with respect to the tiling $\mathbf{T}$. It is easy to see that the exterior

derivative d_k maps $\Lambda^k_{PE}(\mathbf{T})$ to $\Lambda^{k+1}_{PE}(\mathbf{T})$, and we define

$$H^k_{PE}(\mathbf{T},\mathbb{R}) = (\text{Ker } d_k)/(\text{Im } d_{k-1}). \tag{3}$$

Theorem 1.3 ([KP]). *If* $\mathbf{T}$ *is a tiling with finite local complexity with respect to translations, and if* Ω *is the continuous hull of* $\mathbf{T}$*, then* $H^k_{PE}(\mathbf{T},\mathbb{R})$ *is naturally isomorphic to the Čech cohomology of* Ω *with real coefficients, denoted* $\check{H}^k(\Omega,\mathbb{R})$*.*

To get a PE interpretation of integer-valued cohomology, we use the fact that a tiling $\mathbf{T}$ is itself a decomposition of $\mathbb{R}^d$ into 0-cells (vertices), 1-cells (edges). etc. A PE k-cochain α is a function that assigns an integer to each oriented k-cell in a PE way. More precisely, there must be a radius R such that, if c_1 and c_2 are two k-cells with centers of mass x and y, and if $T-x$ and $T-y$ agree on a ball of radius R around the origin, then $\alpha(c_1) = \alpha(c_2)$. (For integer-valued functions, there is no distinction between strong and weak pattern-equivariance.) Let $C^k_{PE}(\mathbf{T})$ denote the set of PE k-cochains. Instead of the exterior derivative, we consider the cellular coboundary map δ_k that maps $C^k_{PE}(\mathbf{T})$ to $C^{k+1}_{PE}(\mathbf{T})$, and define

$$H^k_{PE}(\mathbf{T}) = (\text{Ker } \delta_k)/(\text{Im } \delta_{k-1}). \tag{4}$$

Theorem 1.4 ([Sa2]). *If* $\mathbf{T}$ *is a tiling with finite local complexity with respect to translations, and if* Ω *is the continuous hull of* $\mathbf{T}$*, then* $H^k_{PE}(\mathbf{T})$ *is naturally isomorphic to the Čech cohomology of* Ω *with integer coefficients.*

Sketch of proof. $\mathbf{T}$ induces a map π from $\mathbb{R}^d$ to Ω, sending $x \in \mathbb{R}^d$ to the tiling $\mathbf{T} - x$. Composing with the natural projection from Ω to each approximant Γ^n, we obtain a sequence of maps $\pi_n : \mathbb{R}^d \to \Gamma^n$. The orbit of $\mathbf{T}$ is dense in Ω, so these maps are surjective. Since Γ^n parametrizes the central patch of a tiling, a function on $\mathbb{R}^d$ is (strongly) pattern-equivariant if and only if it is the pullback of a function on one of the approximants Γ^n, and the same goes for cochains. Studying PE cochains of arbitrary radius is equivalent to studying cochains on Γ^n and taking a limit as $n \to \infty$. In other words, $H^k_{PE}(\mathbf{T}) = \varinjlim H^k(\Gamma^n) \simeq \check{H}^k(\Omega)$. □

Example 1. Let $\mathbf{T}$ be a Fibonacci tiling $\ldots babaabaa \ldots$ of $\mathbb{R}$ by long (a) and short (b) tiles. Let i_a be a 1-cochain that evaluates to 1 on each a tile and 0 on each b tile, and let i_b evaluate to 1 on each b and to 0 on each a. Since there are no 2-cells, $\delta i_a = \delta i_b = 0$, so i_a and i_b define classes in $H^1_{PE}(\mathbf{T})$. Once we develop the machinery of Barge–Diamond collaring, we will see that these classes correspond to the generators of $\check{H}^1(\Omega) = \mathbb{Z}^2$.

Example 2. If $\mathbf{T}$ is a Thue–Morse tiling $\ldots abbabaabbaababba \ldots$, obtained from the substitution $a \to ab$, $b \to ba$, one can similarly define indicator 1-cochains i_a and i_b that count a and b tiles. However, these cochains are cohomologous. To see this, divide the tiling $\mathbf{T}$ into 1-supertiles,[4] with each being either ab or ba. Let γ be a PE 0-cochain that evaluates to zero on the vertices that mark the beginning

[4] Recall that if a substitution tiling is non-periodic, then it can be decomposed into supertiles in a unique way, and that this decomposition is a local operation. In the Thue–Morse tiling, every patch of size 5 or greater contains either the sub-word aa or the sub-word bb. The boundaries

or end of such a supertile, to 1 on the vertex in the middle of an ab supertile, and to -1 on the vertex in the middle of a ab supertile. Then $\delta\gamma$ evaluates to 1 on every a tile (since the boundary of an a tile is either the middle vertex of an ab supertile minus the beginning of that supertile, or the end vertex of a ba supertile minus the middle vertex) and -1 on every b tile, so $\delta\gamma = i_a - i_b$.

The first Čech cohomology of the Thue–Morse tiling space is known to be $\mathbb{Z}[1/2] \oplus \mathbb{Z}$. The generators can be chosen as follows. Let α_n be a 1-cochain that evaluates to 1 on the first tile of each n-supertile and to 0 on the other $2^n - 1$ tiles. The cochain α_n basically counts n-supertiles. Since there are two n-supertiles in each $(n+1)$-supertile, α_n is cohomologous to $2\alpha_{n+1}$. Let β be a 1-cochain that evaluates to 1 on each a tile that is followed by a b tile, and to zero on b tiles or on a tiles that are followed by a tiles. This is not cohomologous to any combination of the α_n tiles since, on average, β applied to a long interval yields a third of the length of the interval, something that no finite linear combination of the α_n's can do. The α_n cochains and β generate all of $\check{H}^1$. In this example, the cochains i_a and i_b are both cohomologous to α_1.

1.4. PV cohomology

Another cohomology theory, called PV cohomology after the Pimsner–Voiculscu exact sequence, was developed by Savinien and Bellissard [SB]. This theory is based on the structure of the *transversal* to the tiling space. Since the C^* algebra associated with a tiling space is constructed from the transversal and the associated groupoid, this provides a more intuitive link between the cohomology of a tiling space and the K-theory of the C^* algebra.

We associate a distinguished point, called a *puncture*, to each type of tile. Usually these are chosen in the interior of the tile, say at the center of mass, but the precise choice of puncture is unimportant. The *canonical transversal* Ξ of a tiling space is the set of tilings for which there is a puncture at the origin. This is a Cantor set, and we can study the ring of continuous integer-valued functions on Ξ, denoted $C(\Xi, \mathbb{Z})$. If α is a d-cochain, we define an associated function f_α on Ξ as follows: if $\mathbf{T} \in \Xi$, then $f_\alpha(\mathbf{T})$ equals α applied to the tile of $\mathbf{T}$ that lies at the origin. This map induces an isomorphism (as an additive group) between $C(\Xi, \mathbb{Z})$ and $C^d_{PE}(\mathbf{T})$.

Similarly, we can define punctures for all of the lower-dimensional faces and edges and vertices of different tiles, with the condition that if (say) an edge is on the boundary of two tiles, then its puncture viewed as the boundary of the first tile is the same as its puncture viewed as the boundary of the second tile. For k ranging from 0 to d, let Ξ^k_Δ be the set of tilings where the origin sits at a puncture of an k-cell. As with $\Xi = \Xi^d_\Delta$, $C(\Xi^k_\Delta, \mathbb{Z})$ is isomorphic to $C^k_{PE}(\mathbf{T})$.

In PV cohomology, the group of k-cochains is $C(\Xi^k_\Delta, \mathbb{Z})$ and the coboundary maps are built from the geometry of the specific tiles. After untangling the definitions, these coboundary maps turn out to be identical to the coboundary maps

between 1-supertiles sit in the middle of these sub-words, and at all points at even distance from these middles.

in PE-cohomology. Thus, PV theory and PE theory not only have the same cohomologies, but have isomorphic cochain complexes. For details of this argument, see [BK].

1.5. Quotient cohomology

So far we have been discussing the absolute cohomology of each tiling space. However, cohomology is also a functor that concerns maps between spaces. Inclusions give rise to relative cohomology (see [Hat]), while surjections give rise to a less-known construction called *quotient cohomology*.

Let $f : \Omega_X \to \Omega_Y$ be a factor map of tiling spaces. As long as the tilings have finite local complexity with respect to translations, the pullback map f^* is an injection on cochains. (This argument applies both to Čech cochains and to pattern-equivariant cochains.) We then define the quotient cochain complex $C^k_Q(\Omega_X, \Omega_Y)$ to be $C^k(\Omega_X)/f^*(C^k(\Omega_Y))$, and the quotient cohomology $H^k_Q(\Omega_X, \Omega_Y)$ to be the cohomology of this complex. The short exact sequence of cochain complexes:

$$0 \to C^k(\Omega_Y) \xrightarrow{f^*} C^k(\Omega_X) \to C^k_Q(\Omega_X, \Omega_Y) \to 0 \tag{5}$$

induces a long exact sequence of cohomology groups

$$\cdots \to \check{H}^k(\Omega_Y) \xrightarrow{f^*} \check{H}^k(\Omega_X) \to H^k_Q(\Omega_X, \Omega_Y) \to \check{H}^{k+1}(\Omega_Y) \to \cdots . \tag{6}$$

As with ordinary relative (co)homology, there is an excision principle:

Theorem 1.5 ([BSa]). *Let $f : X \to Y$ be a quotient map such that f^* is an injection on cochains. If $Z \subset X$ is an open set such that f is injective on the closure of Z, then $H^k_Q(X, Y)$ is isomorphic to $H^k_Q(X - Z, Y - f(Z))$.*

For factor maps between tiling spaces, excision cannot be used directly. Every orbit is dense, so there are no open sets where f is injective on the closure. However, it is often the case that a factor map $\Omega_X \to \Omega_Y$ is injective apart from a small set of tilings. In such circumstances, one can use homotopy to convert the tiling spaces into spaces where excision does apply.

Example 3. Let Ω_X be the one-dimensional tiling space obtained from the period-doubling substitution $a \to ab, b \to aa$, and let Ω_Y be the dyadic solenoid Sol_2 which can be viewed formally as coming from a substitution $c \to cc$. (The dyadic solenoid is not actually a tiling space, but it has similar topological properties, being an inverse limit space, allowing us to apply the machinery of quotient cohomology.) There is a factor map $f : \Omega_X \to \Omega_Y$ that identifies two translational orbits but is otherwise injective. This map sends a tiling $\mathbf{T}$ to the sequence $(x_0, x_1, \ldots)$, where x_k is the location of the endpoints of the k-supertiles (mod 2^k). In other words, $f(\mathbf{T})$ gives the locations of the supertiles of all order in $\mathbf{T}$, but does not indicate which supertiles are of type a or type b. However, in a period-doubling tiling the nth order supertiles are identical except on the very last entry. Unless the tiling $\mathbf{T}$ consists of two infinite-order supertiles, $f(\mathbf{T})$ determines $\mathbf{T}$. If the tiling $\mathbf{T}$ does

consist of two infinite-order supertiles, then there is exactly one other tiling $\mathbf{T}'$, differing from $\mathbf{T}$ only at a single letter, such that $f(\mathbf{T}') = f(\mathbf{T})$.

In that last instance, we say that $\mathbf{T}$ and $\mathbf{T}'$ have a *zero-dimensional feature*, namely the boundary of an infinite-order supertile, and agree away from that feature. Let $\Omega_{X_0} = \{\mathbf{T}, \mathbf{T}'\}$, and let $\Omega_{Y_0} = f(\mathbf{T})$. The map f is basically a quotient map, identifying the orbit of $\mathbf{T}$ with the orbit of $\mathbf{T}'$. This identification is the *suspension* of a map from the 2-point set Ω_{X_0} to the 1-point set Ω_{Y_0}.

The situation of this example is quite common. There are many situations where a factor map $f : \Omega_X \to \Omega_Y$ between tiling spaces (or solenoids) is injective except on the translational orbits of a set Ω_{X_0} of tilings. Furthermore, Ω_{X_0} has the structure of a $d - \ell$-dimensional tiling space, admitting an $\mathbb{R}^{d-\ell}$ action and locally being the product of $\mathbb{R}^{d-\ell}$ and a totally disconnected set. Defining Ω_{Y_0} to be $f(\Omega_{X_0})$, the following theorem relates the quotient cohomologies of (Ω_X, Ω_Y) and $(\Omega_{X_0}, \Omega_{Y_0})$.

Theorem 1.6 ([BSa]). *Let $f : \Omega_X \to \Omega_Y$ be a quotient map of tiling spaces such that f^* is injective on cochains. Suppose that f is injective aside from the translational orbits of a codimension-ℓ set $\Omega_{X_0} \subset \Omega_X$ of tilings. Let $\Omega_{Y_0} = f(\Omega_{X_0})$. Then $H^k_Q(\Omega_X, \Omega_Y) = H^{k-\ell}_Q(\Omega_{X_0}, \Omega_{Y_0})$.*

In our example, $\ell = 1$, Ω_{X_0} consists of two points, Ω_{Y_0} is a single point, $H^0_Q(\Omega_{X_0}, \Omega_{Y_0}) = \mathbb{Z}$, and so $H^1_Q(\Omega_X, \Omega_Y) = \mathbb{Z}$. Since $\check{H}^1(\Omega_Y) = \mathbb{Z}[1/2]$, the long exact sequence (6) shows that $\check{H}^1(\Omega_X) = \mathbb{Z}[1/2] \oplus \mathbb{Z}$. This is in fact the first cohomology of the period-doubling space.

An extension of Theorem 1.6 relates the generators of $H^{k-\ell}_Q(\Omega_{X_0}, \Omega_{Y_0})$ to the generators of $H^k_Q(\Omega_X, \Omega_Y)$. This allows us to construct generators for $H^k(\Omega_X)$ from generators of $H^k(\Omega_Y)$ and from generators of $H^*_Q(\Omega_{X_0}, \Omega_{Y_0})$.

2. How do you compute tiling cohomology?

As with other topological spaces, there is no single "best" method for computing the cohomology of a tiling space. Different tiling spaces are best addressed with different methods.

Cut-and-project tiling spaces are measurably conjugate to Kronecker flows on higher-dimensional tori. As topological spaces, they are obtained from the tori by removing some hyperplanes and gluing them back in multiple times. Forrest, Hunton and Kellendonk [FHK], and later Kalugin [Kal] developed ways to compute the cohomology of Ω from the geometry of the "window" used in the cut-and-project scheme.

Substitution tilings can easily be expressed as inverse limits spaces in which all the approximants Γ^n are homeomorphic to a single space Γ^0, and where the substitution σ can be viewed as a map from Γ^0 to itself. For these spaces, computing the cohomology boils down to understanding the cohomology of Γ^0 and

tracking how the classes evolve under the pullback map σ^*. There are many ways to do this, and each inverse limit scheme gives rise to a calculational method. In this section we develop several such schemes, beginning with the original ideas of Anderson and Putnam, and working our way through Gähler's construction and the more recent ideas of Barge, Diamond, Hunton and Sadun. Variants of the Anderson–Putnam and Barge–Diamond methods are then applied to tilings with rotational symmetry and to hierarchical tilings that are not substitutions (e.g., the "generalized substitutions" of [F, AFHI]).

Tilings that come from local matching rules are harder to understand. However, they can sometimes be related to substitution tilings [Moz, GS, Rad]. When a substitution tiling space Ω_Y is the quotient of a local matching rules tiling space Ω_X, we can study the cohomology of Ω_X via the cohomology of Ω_Y and the quotient cohomology $H^k_Q(\Omega_X, \Omega_Y)$.

2.1. The Anderson–Putnam complex

Suppose that we have a substitution tiling whose tiles are polygons that meet full-edge to full-edge. We construct an inverse limit space whose approximants Γ^n describe partial tilings. Specifically, a point in Γ^n describes where the origin sits within an n-supertile. Since this also determines where the origin sits within an $(n-1)$-supertile, we have a natural map $\sigma : \Gamma^n \to \Gamma^{n-1}$ and can consider the inverse limit space $\Omega^0 = \varprojlim(\Gamma^n, \sigma)$.

Since the origin can sit anywhere in a supertile of any type, Γ^n consists of one copy of each type of supertile. However, there is an ambiguity when the origin sits on the boundary of a supertile. If the origin sits on the boundary between supertile A and supertile B, do we consider it as part of A or B? The answer is to identify the two edges.

Specifically, Γ^n is obtained by taking the disjoint union of one copy of each kind of (closed) n-supertile, and then applying the relation that, if somewhere in an admissible tiling an edge e_1 of supertile A coincides with an edge e_2 of supertile B, then e_1 and e_2 are identified.

These identifications do not just come in pairs. It may happen that the right edge of A is identified with the left edges of both B and C, and that the left edge of C is identified with the right edges of both A and D. In that case, the left edges of B and C and the right edges of A and D would all be identified. The information contained in that point in Γ^n would indicate that the origin either sits at a particular spot on the right edge of A, or at that spot on the right edge of D, and also that it sits at the corresponding spot on the left edge of either B or C.

The set of possible n-supertiles looks just like the set of possible tiles, only scaled up by a factor of λ^n. As a result, Γ^n is just a scaled-up version of Γ^0. Γ^0 is called the (uncollared) Anderson–Putnam complex of the substitution σ, and is denoted Γ_{AP}. Furthermore, the decomposition of n-supertiles into constituent $(n-1)$-supertiles is combinatorially the same for all n. After rescaling, there is a single map (which we again call σ) from Γ_{AP} to itself. This map involves stretching

each tile in Γ_{AP} by a factor of λ, dividing it into tiles via the substitution rule, and then identifying pieces. We then define $\Omega^0 = \varprojlim(\Gamma_{AP}, \sigma)$.

2.1.1. Forcing the border. *Forcing the border* was defined by Johannes Kellendonk in his study [Kel1] of the Penrose tiling. As we shall see, if a substitution forces the border, then Ω^0 is homeomorphic to the tiling space Ω, allowing for an easy computation of the cohomology of Ω. If a substitution doesn't force the border, then there are a variety of *collaring* techniques for describing the tiling space via a slightly different substitution that does force the border. By combining collaring with the Anderson–Putnam construction, we can compute the cohomology of arbitrary substitution tiling spaces.

Suppose we have a non-periodic substitution tiling space, so that $\sigma : \Omega \to \Omega$ is a homeomorphism [Mos, Sol]. This means that we can decompose each tiling T uniquely into a collection of non-overlapping 1-supertiles, and by extension we can decompose T uniquely into non-overlapping k-supertiles for every k. The substitution is said to *force the border at level k* if any two k-supertiles of the same type not only have the same decomposition into tiles, but also have the same pattern of ordinary tiles surrounding them (i.e., the pattern of tiles that touch the supertiles at 1 or more points). Moreover, any two n-supertiles with $n > k$ have the same pattern of $(n - k)$-supertiles surrounding them.

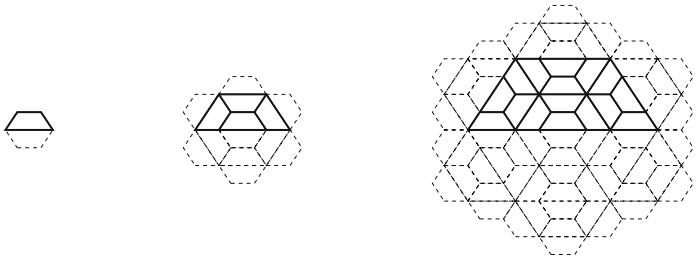

FIGURE 1. In bold face, a half-hex tile, an order-1 supertile, and an order-2 supertile. In dotted lines, the nearby tiles that these determine.

The half-hex substitution is shown in Figure 1. The solid lines indicate the tiles within a supertile, and the dotted lines indicate the neighboring tiles that must also appear. This substitution forces the border at level 2, since the 2-supertile is completely surrounded by determined tiles, but does not force the border at level 1, since some of the tiles that touch the four vertices of the 1-supertile are undetermined. By contrast, the chair tiling does not force the border at all, since tiles near the southwest corner of a chair supertile of arbitrary order can appear in either of the patterns shown in Figure 2.

If a substitution forces the border at level k, then a point in Γ^n not only determines where the origin sits in a supertile of level n, but it determines all of the $(n - k)$-supertiles surrounding the supertile that contains the origin. If the origin sits on the boundary between two or more n-supertiles, then there is some

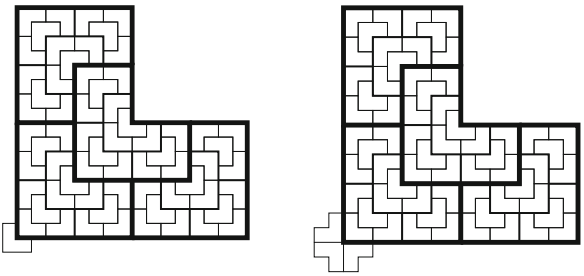

FIGURE 2. There are two ways to extend a high-order chair supertile around the southwest corner.

ambiguity on the nature of the n-supertiles that surround the origin. However, there is no ambiguity about the $(n-k)$-supertiles that surround the origin.

The inverse limit $\Omega^0 = \varprojlim(\Gamma^n_{AP}, \sigma)$ is a sequence of consistent instructions for placing higher and higher-order supertiles in a growing region containing the origin. The union of these regions is all of $\mathbb{R}^d$. This is tantamount to

Theorem 2.1. *If σ is a substitution that forces the border and has finite local complexity with respect to translations, then the corresponding tiling space Ω is homeomorphic to Ω^0.*

2.1.2. Anderson–Putnam collaring. If the substitution σ does not force the border, then Ω^0 is typically not homeomorphic to Ω. There is still a map $\Omega \to \Omega^0$, whose nth coordinate is a description of the n-supertile containing the origin. Furthermore, this map is surjective. However, it is typically not injective. Even if the origin is not on a boundary, knowing the supertiles to all orders containing the origin may not describe the entire tiling, since the union of these supertiles may be a quarter-plane or a half-plane. If there is more than one extention of this infinite partial-tiling to the entire plane, then there is more than one preimage in Ω.

To remedy this, we construct a new substitution using collared tiles. Take a tiling $\mathbf{T}$, and identify tiles that are (a) of the same type and (b) whose nearest neighbors are all of the same type. That is, tiles t_1 and t_2 are identified if, for some points $x \in t_1$ and $y \in t_2$, the tilings $\mathbf{T} - x$ and $\mathbf{T} - y$ agree exactly on the tile containing the origin and on all tiles touching this central tile. A *collared tile* is an equivalence class of tiles under this identification. Note that a collared tile has the *same size and shape* as an ordinary uncollared tile. The difference is that the label of the collared tile carries extra information about its surroundings.

Example 4. In the Fibonacci tiling, every b tile is preceded and followed by an a tile, while an a tile has three possibilities for its neighbors. There are thus four collared tiles, which we denote $A_1 = (a)a(b)$, $A_2 = (b)a(a)$, $A_3 = (b)a(b)$ and $B = (a)b(a)$, where the notation $(x)y(z)$ means a y tile that is preceded by an x and followed by a z. Under substitution, $A_1 \to (ab)ab(a) = A_3B$, $A_2 \to (a)ab(ab) = A_1B$, $A_3 \to (a)ab(a) = A_1B$, and $B \to (ab)a(ab) = A_2$.

We can relabel all of our tiles according to their neighbors to obtain a new tiling by collared tiles. For instance, in the Fibonacci tiling the pattern $\ldots babaabaababaa\ldots$ becomes $\ldots BA_3BA_2A_1BA_2A_1BA_3BA_2A_1\ldots$.

Theorem 2.2 ([AP]). *Rewriting a substitution in terms of collared tiles always yields a system that forces the border.*

Sketch of proof. A collared tile is a tile together with a pattern of nearest neighbors, thereby determining all the tiles in at least an ϵ-neighborhood. After substituting n times, we obtain an n-supertile together with a pattern of neighboring n-supertiles, thereby determining all the tiles within a distance $\lambda^n\epsilon$. Pick n big enough that $\lambda^n\epsilon$ is more than twice the diameter of the largest tile. The n-times substituted (collared) tile then determines its neighboring uncollared tiles *and the neighbors of these neighbors*, and hence determines its neighboring collared tiles. □

For instance, in the Fibonacci example,

$$\begin{aligned}
\sigma^2(A_1) &= (aba)aba(ab) = (BA_2)A_1BA_2(A_1),\\
\sigma^2(A_2) &= (ab)aba(aba) = (B)A_3BA_2(A_1B),\\
\sigma^2(A_3) &= (ab)aba(ab) \;\; = (B)A_3BA_2(A_1),\\
\sigma^2(B) &= (aba)ab(aba) = (BA_2)A_1B(A_3B).
\end{aligned} \tag{7}$$

In each case, substituting a collared tile twice determines at least two extra tiles on each side of the 2-supertile, and so determines the collared tile on each side of the supertile. Combining this theorem with the first Anderson–Putnam construction yields the following

Theorem 2.3 ([AP]). *Let Ω be a tiling space derived from a substitution σ. Assume that there are only finitely many tile types, up to translation, and that the tiles are polygons (or polyhedra) that meet full edge to full edge (or full face to full face). Then Ω is homeomorphic to $\varprojlim(\tilde{\Gamma}_{AP}, \sigma)$, where $\tilde{\Gamma}_{AP}$ is constructed using once-collared tiles.*

2.2. Gähler's construction

One can iterate the collaring construction, rewriting an arbitrary tiling space Ω in terms of collared tiles, then in terms of collared collared tiles (i.e., two tiles of the same type are identified only if they have the same pattern of nearest and second-nearest neighbors), and more generally n-times collared tiles. Let Γ_G^n be the Anderson–Putnam complex constructed from the n-times collared tiles. There is a natural quotient map $q_n : \Gamma_G^n \to \Gamma_G^{n-1}$ that merely forgets about the nth nearest neighbors.

Theorem 2.4. *Let Ω be any space of tilings that have finite local complexity with respect to translation. Then Ω is homeomorphic to the inverse limit of the approximants Γ_G^n under the forgetful maps q_n.*

Sketch of proof. (see [Gah, Sa1]) A point $p_n \in \Gamma_G^n$ is either a point in an n-collared tile, or is the identification of several possible points on the boundary of an n-collared tile. Either way, at least $n-1$ rings of tiles around p_n are specified. The point p_n can then be viewed as instructions for building a patch *around the origin*. A sequence $p_0, p_1, \ldots$ is then a consistent set of instructions for building larger and larger patches around the origin, whose union is $\mathbb{R}^d$. Hence $\varprojlim(\Gamma_G^n, q_n)$ parametrizes tilings in Ω. □

Gähler's construction is extremely useful for theoretical arguments, as it applies to all tiling spaces, not just to substitution tiling spaces. For instance, the identification of integer-valued pattern-equivariant cohomology with Čech cohomology [Sa2] is based on this construction. Unfortunately, it has not proven effective in computing cohomology. $\check{H}^*(\Omega)$ does equal $\varinjlim H^*(\Gamma_G^n)$, but there is no general procedure for computing $H^*(\Gamma_G^n)$. The number of cells in Γ_G^n grows with n, and it is difficult to do computations that apply simultaneously to all values of n.

2.3. Barge–Diamond collaring

The Anderson–Putnam and Gähler constructions are based on collared *tiles*. The Barge–Diamond construction [BD2, BDHS] is based on collared *points*.

Let $\mathbf{T} \in \Omega$ be a non-periodic substitution tiling. Recall that non-periodicity implies that the substitution σ has an inverse on Ω. Pick a radius r and consider the equivalence relation on $\mathbb{R}^d$: $x \sim y$ if the tilings $\mathbf{T} - x$ and $\mathbf{T} - y$ agree out to distance r around the origin. Likewise, let $x \sim_n y$ if the tilings $\sigma^{-n}(\mathbf{T} - x)$ and $\sigma^{-n}(\mathbf{T} - y)$ agree out to distance r. That is, if $\mathbf{T} - x$ and $\mathbf{T} - y$ have the same structure of n-supertiles out to distance $\lambda^n r$. (In particular, they also have the same structure of ordinary tiles out to distance $\lambda^n r$.) Let Γ_{BD}^n be the quotient of $\mathbb{R}^d$ by $\sim_n$. A priori this would seem to depend on the tiling $\mathbf{T}$, but for minimal tiling spaces all tilings have the same patterns and give rise to identical approximants. Since $x \sim_n y$ implies $x \sim_{n-1} y$, there is a natural quotient map $q_n : \Gamma_{BD}^n \to \Gamma_{BD}^{n-1}$. Furthermore, the complexes Γ_{BD}^n are all homeomorphic. Indeed, if $\mathbf{T}$ is a self-similar tiling with $\sigma(\mathbf{T}) = \mathbf{T}$, then $x \sim y$ if and only if $\lambda^n x \sim_n \lambda^n y$, so Γ_{BD}^n is just an enlarged copy of a single space Γ_{BD} and the quotient maps q_n are all induced from the substitution σ.

The radius r is arbitrary, but for many applications it is convenient to take r extremely small. The complex Γ_{BD} is then a CW complex comprised of pieces of tiles. For instance, suppose that $\mathbf{T}$ is a one-dimensional tiling. Points x and y are identified if either (1) they are in corresponding places in tiles of the same type, and are farther than r from the nearest vertex, or (2) they are in corresponding places in tiles of the same type, within distance r of a vertex, and the tiles on the other side of the vertices are the same. If the tiles all have length 1, then the equivalence classes of the first type form 1-cells of length $1 - 2r$, one for each tile type. We call these *tile cells*. The equivalence classes of the second type form 1-cells of length $2r$, called *vertex flaps*, one for each possible transition from one tile to another. For instance, in the Fibonacci tiling, the possible 2-tile patches are aa, ab,

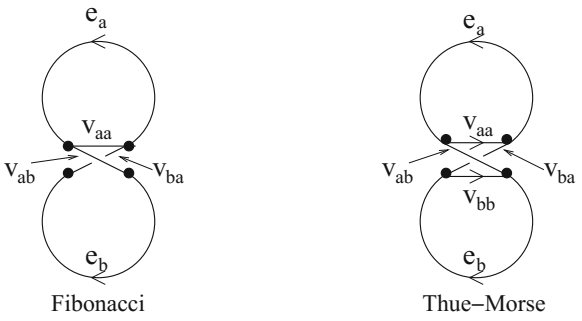

FIGURE 3. Barge–Diamond Complexes for the Fibonacci and Thue–Morse Substitutions

and ba, so Γ_{BD} consists of two tile cells (a and b) and three vertex flaps, arranged as in Figure 3. In the Thue–Morse tiling, all four transitions $\{aa, ab, ba, bb\}$ are possible, so we have two edge cells and four vertex flaps, also shown in Figure 3.

In a two-dimensional tiling, there are three kinds of 2-cells. *Tile cells* correspond to the interiors of tiles, *edge flaps* correspond to points that are within r of an edge, and contain information about what tile is on the other side of the edge, and *vertex polygons* describe what is happening near a vertex, and have information about all of the tiles touching the vertex. If the tiles are unit squares meeting edge-to-edge, then the tile cells are $(1-2r)\times(1-2r)$ squares, the edge flaps are $2r\times(1-2r)$ rectangles, and the vertex polygons are $2r\times 2r$ squares. (Strictly speaking, this requires using the L^∞ metric on $\mathbb{R}^2$ rather than the Euclidean metric, to avoid having arcs of circles on the boundaries of cells.)

Theorem 2.5 ([BD2, BDHS]). *For any positive radius r, Ω is homeomorphic to the inverse limit $\varprojlim(\Gamma_{BD}, \sigma)$.*

Proof. As with the Anderson–Putnam construction, a point in the inverse limit is a sequence of instructions for tiling larger and larger regions of the plane, insofar as the nth approximant determines the structure of a tiling out to distance $\lambda^n r$. □

The complexes Γ^n_{BD} are all the same (up to scale), so it is relatively easy to compute $H^*(\Gamma^n_{BD}) = H^*(\Gamma_{BD})$. Unfortunately, the map $\sigma : \Gamma_{BD} \to \Gamma_{BD}$ is typically not a cellular map. For instance, for a square tiling σ takes a $2r\times 2r$ vertex polygon to a $2\lambda r\times 2\lambda r$ square, which is a vertex polygon plus a small piece of the adjacent edge flaps and tile cells. To do our computations we need to use a map $\tilde\sigma$ that is cellular and homotopic to σ. (One way to get such a map $\tilde\sigma$ is to compose σ with a flow that expands tile cells slightly at the expense of the edge cells and vertex polygons. The details are *not* important.) The map $\tilde\sigma$ sends vertex polygons to vertex polygons, edge flaps to a union of edge flaps and vertex polygons, and tile cells to a union of all three kinds of cells. Let $\tilde\Omega = \varprojlim(\Gamma_{BD}, \tilde\sigma)$.

Theorem 2.6. *The Čech cohomology of $\tilde{\Omega}$ is isomorphic to the Čech cohomology of Ω.*

Proof. Since σ and $\tilde{\sigma}$ are homotopic, $\tilde{\sigma}^* = \sigma^*$ as operators on $H^*(\Gamma_{BD})$. Then

$$\begin{aligned} \check{H}^*(\tilde{\Omega}) &= \check{H}^*(\varprojlim(\Gamma_{BD}, \tilde{\sigma})) = \varinjlim H^*(\Gamma_{BD}, \tilde{\sigma}^*) = \varinjlim H^*(\Gamma_{BD}, \sigma^*) \qquad (8) \\ &= \check{H}^*(\varprojlim(\Gamma_{BD}, \sigma)) = \check{H}^*(\Omega). \end{aligned}$$

□

This theorem does *not* say that $\tilde{\Omega}$ and Ω are homeomorphic. In many cases they are not. However, their cohomologies are the same, so we can always use the inverse limit structure of $\tilde{\Omega}$ to compute the cohomology of Ω.

2.3.1. One-dimensional results ([BD2]). Let $S_0 \subset \Gamma_{BD}$ be the sub-complex of vertex flaps, and let $S_1 = \Gamma_{BD}$. Since $\tilde{\sigma}$ maps S_0 to S_0 and S_1 to S_1, we can consider the inverse limit space $\mathcal{S}_i = \varprojlim(S_i, \tilde{\sigma})$. Since $\mathcal{S}_0 \subset \mathcal{S}_1$, we can compute $\check{H}^*(\Omega) = \check{H}^*(\mathcal{S}_1)$ by computing $\check{H}^*(\mathcal{S}_0)$ and the relative cohomology $\check{H}^*(\mathcal{S}_1, \mathcal{S}_0)$ and then combining them with the long exact sequence

$$0 \to \check{H}^0(\mathcal{S}_1, \mathcal{S}_0) \to \check{H}^0(\mathcal{S}_1) \to \check{H}^0(\mathcal{S}_0) \to \check{H}^1(\mathcal{S}_1, \mathcal{S}_0) \to \check{H}^1(\mathcal{S}_1) \to \check{H}^1(\mathcal{S}_0) \to 0 \qquad (9)$$

We examine each of these terms. $\check{H}^0(\mathcal{S}_1, \mathcal{S}_0)$ is the direct limit (under $\tilde{\sigma}^*$) of $H^0(S_1, S_0)$. Since S_1 is connected, this is zero. Likewise, $\check{H}^0(\mathcal{S}_1) = \varinjlim H^0(\mathcal{S}_1) = \mathbb{Z}$. Since $\tilde{\sigma}$ maps each cell of S_0 to a single cell, $\tilde{\sigma}$ merely permutes the cells of the eventual range S_0^{ER}. Thus $\varinjlim H^*(S_0) = H^*(S_0^{ER})$. If S_0^{ER} has k connected components and has ℓ loops, then $\check{H}^0(\mathcal{S}_0) = \mathbb{Z}^k$ and $\check{H}^1(\mathcal{S}_0) = \mathbb{Z}^\ell$. Meanwhile the quotient space S_1/S_0 is a wedge of circles, one for each tile type. $H^1(S_1, S_0) = \mathbb{Z}^N$, where N is the number of tile types, and $\check{H}^1(\mathcal{S}_1, \mathcal{S}_0) = \varinjlim(\mathbb{Z}^N, A^T)$, where A is the substitution matrix. Combining these observations, we have the long exact sequence

$$0 \to \mathbb{Z} \to \mathbb{Z}^k \to \varinjlim(\mathbb{Z}^N, A^T) \to \check{H}^1(\Omega) \to \mathbb{Z}^\ell \to 0. \qquad (10)$$

Using reduced cohomology, this can be further simplified to

$$0 \to \mathbb{Z}^{k-1} \to \varinjlim(\mathbb{Z}^N, A^T) \to \check{H}^1(\Omega) \to \mathbb{Z}^\ell \to 0. \qquad (11)$$

In the Fibonacci tiling, $A = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and S_0 consists of three vertex flaps: aa, ab, and ba. These form a contractible set, so $k - 1 = \ell = 0$, and $\check{H}^1(\Omega) = \varinjlim(\mathbb{Z}^2, A^T) = \mathbb{Z}^2$. In fact, whenever S_0^{ER} is contractible, $\tilde{H}^0(S_0^{ER})$ and $H^1(S_0^{ER})$ vanish and $H^1(\Omega)$ is isomorphic to $\varinjlim(\mathbb{Z}^N, A^T)$.

We can also describe the Fibonacci tiling using collared tiles $A_1 = (a)a(b)$, $A_2 = (b)a(a)$, $A_3 = (b)a(b)$, and $B = (a)b(a)$. Collaring the Fibonacci tiles and then applying the Barge–Diamond construction is overkill, but this example shows the interplay of the substitution matrix and the cohomology of S_0^{ER}. Our complex Γ_{BD} has four tile cells and five vertex flaps, namely A_1B, A_2A_1, A_3B, BA_2, and BA_3. However, A_1B and A_3B are not in S_0^{ER}, since all supertiles start with A_1, A_2, or A_3. S_0^{ER} consists of just the flaps A_2A_1, BA_2 and BA_3, yielding $k = 2$

and $\ell = 0$. The substitution matrix is $\begin{pmatrix} 0&1&1&0\\ 0&0&0&1\\ 1&0&0&0\\ 1&1&1&0 \end{pmatrix}$. This matrix has rank 3, with eigenvalues $(1 \pm \sqrt{5})/2$, -1, and 0, and $\varinjlim(\mathbb{Z}^4, A^T) = \mathbb{Z}^3$. We then have $0 \to \mathbb{Z} \to \mathbb{Z}^3 \to \check{H}^1(\Omega) \to 0$. After checking that the quotient of $\mathbb{Z}^3$ by $\mathbb{Z}$ is $\mathbb{Z}^2$ (with no torsion terms), we again obtain $\check{H}^1(\Omega) = \mathbb{Z}^2$.

In the Thue–Morse tiling, $M = \left(\begin{smallmatrix} 1&1\\ 1&1 \end{smallmatrix}\right)$ and S_0 consists of four vertex flaps that form a loop. Now $\varinjlim(\mathbb{Z}^2, A^T) = \mathbb{Z}[1/2]$ and $k = \ell = 1$, so we have $0 \to \mathbb{Z}[1/2] \to \check{H}^1(\Omega) \to \mathbb{Z} \to 0$. Since $\mathbb{Z}$ is free, this sequence splits, so $\check{H}^1(\Omega) = \mathbb{Z}[1/2] \oplus \mathbb{Z}$.

2.3.2. Higher dimensions ([BDHS]). In higher dimensions the procedure is similar, but the results cannot be expressed in a single exact sequence such as (11). In two dimensions, we consider the complex S_0 of vertex polygons, S_1 of vertex polygons and edge flaps, and $S_2 = \Gamma_{BD}$. We also consider the inverse limits $\mathcal{S}_i = \varprojlim(S_i, \tilde{\sigma})$. As in one dimension, $\tilde{\sigma}$ maps each vertex polygon to a single vertex polygon, so $\check{H}^*(\mathcal{S}_0) = H^*(S_0^{ER})$. However, S_0 is a two-dimensional complex, so computing the cohomology of S_0^{ER} is more than just counting connected components and loops.

The next step is to consider $\check{H}^*(\mathcal{S}_1, \mathcal{S}_0) = \varinjlim(\check{H}^*(S_1/S_0), \tilde{\sigma}^*)$. This involves only the eventual range of S_1, but is typically a complicated calculation. The quotient space S_1/S_0 breaks into several pieces, one for each direction that an edge can point. In general, the pieces are not particularly simple, and it takes work to understand how $\tilde{\sigma}^*$ acts on $\tilde{H}^*(S_1/S_0)$. Once $\check{H}^*(\mathcal{S}_1, \mathcal{S}_0)$ is computed, we combine it with $\check{H}^*(\mathcal{S}_0)$ via the long exact sequence

$$\cdots \to \check{H}^k(\mathcal{S}_1, \mathcal{S}_0) \to \check{H}^k(\mathcal{S}_1) \to \check{H}^k(\mathcal{S}_0) \to \check{H}^{k+1}(\mathcal{S}_1, \mathcal{S}_0) \to \cdots \tag{12}$$

to compute $\check{H}^*(\mathcal{S}_1)$.

The relative cohomology $\check{H}^*(\mathcal{S}_2, \mathcal{S}_1)$ is simpler. The quotient space S_2/S_1 is a wedge of spheres, so $\tilde{H}^0 = \tilde{H}^1 = 0$ and $\tilde{H}^2 = \mathbb{Z}^N$. $\check{H}^k(\mathcal{S}_2, \mathcal{S}_1)$ equals $\varinjlim(\mathbb{Z}^n, A^T)$ when $k = 2$, and vanishes when $k = 0$ or 1. The final stage is combining $\check{H}^*(\mathcal{S}_1)$ and $\check{H}^*(\mathcal{S}_2, \mathcal{S}_1)$ with the long exact sequence

$$\cdots \to \check{H}^k(\mathcal{S}_2, \mathcal{S}_1) \to \check{H}^k(\mathcal{S}_2) \to \check{H}^k(\mathcal{S}_1) \to \check{H}^{k+1}(\mathcal{S}_2, \mathcal{S}_1) \to \cdots . \tag{13}$$

Example 5. Consider a tiling of $\mathbb{R}^2$ featuring three square tiles A, B, and C, and generated by the substitution $\boxed{*} \to \begin{array}{|c|c|}\hline A & * \\ \hline B & C \\ \hline \end{array}$, where "$*$" is shorthand for A, B or C. This substitution does not force the border, so collaring is needed to compute its cohomology. S_0 involves many vertex polygons, but each of these maps to a vertex polygon of the form $\begin{array}{|c|c|}\hline C & B \\ \hline * & A \\ \hline \end{array}$. S_0^{ER} is contractible, consisting of three squares glued together at their north and east edges, so $\check{H}^0(\mathcal{S}_0) = \mathbb{Z}$ and $\check{H}^1(\mathcal{S}_0) = \check{H}^2(\mathcal{S}_0) = 0$.

S_1/S_0 consists of vertical and horizontal edge flaps. The vertical edge flaps are $B|C$, $C|B$, $A|*$ and $*|A$, but only $C|B$ and $*|A$ survive to the eventual range. This portion of S_1^{ER}/S_0 retracts to the wedge of two circles, and $\tilde{\sigma}^*$ acts on its first

cohomology by the matrix $\left(\begin{smallmatrix}1&1\\1&1\end{smallmatrix}\right)$, yielding a direct limit of $\mathbb{Z}[1/2]$. The horizontal edge flaps are similar, giving another factor of $\mathbb{Z}[1/2]$, so $\check{H}^1(\mathcal{S}_1, \mathcal{S}_0) = \mathbb{Z}[1/2]^2$ and $\check{H}^0(\mathcal{S}_1, \mathcal{S}_0) = \check{H}^2(\mathcal{S}_1, \mathcal{S}_0) = 0$.

S_2/S_0 is a wedge of three spheres, and the only nontrivial cohomology is $H^2 = \mathbb{Z}^3$. This transforms via $A^T = \left(\begin{smallmatrix}2&1&1\\1&2&1\\1&1&2\end{smallmatrix}\right)$, so $\check{H}^2(\mathcal{S}_2, \mathcal{S}_1) = \varinjlim(\mathbb{Z}^3, A^T)$ and $\check{H}^0(\mathcal{S}_2, \mathcal{S}_1) = \check{H}^1(\mathcal{S}_2, \mathcal{S}_1) = 0$.

We combine these relative cohomologies using the long exact sequences (12) and (13). The first of these yields:

$$0 \to 0 \to \check{H}^0(\mathcal{S}_1) \to \mathbb{Z} \to \mathbb{Z}[1/2]^2 \to \check{H}^1(\mathcal{S}_1) \to 0, \tag{14}$$

so $\check{H}^0(\mathcal{S}_1) = \mathbb{Z}$ and $\check{H}^1(\mathcal{S}_1) = \mathbb{Z}[1/2]^2$ (and $\check{H}^2(\mathcal{S}_1) = 0$). The second yields:

$$0 \to \check{H}^1(\mathcal{S}_2) \to \mathbb{Z}[1/2]^2 \to \varinjlim(\mathbb{Z}^3, A^T) \to \check{H}^2(\mathcal{S}_2) \to 0. \tag{15}$$

All maps commute with $\tilde{\sigma}^*$. Since the $\mathbb{Z}[1/2]^2$ terms double with substitution, and since the eigenvalues of A^T are 1, 1, and 4, the map from $\mathbb{Z}[1/2]^2$ to $\varinjlim(\mathbb{R}^3, A^T)$ must be zero. We then have

$$\check{H}^0(\mathcal{S}_2) = \mathbb{Z}, \qquad \check{H}^1(\mathcal{S}_2) = \mathbb{Z}[1/2]^2, \qquad \check{H}^2(\mathcal{S}_2) = \varinjlim(\mathbb{Z}^3, A^T) = \mathbb{Z}[1/4] \oplus \mathbb{Z}^2. \tag{16}$$

(We write $\mathbb{Z}[1/4]$ rather than $\mathbb{Z}[1/2]$ in $\check{H}^2$ to emphasize that this term scales by 4 under substitution.) This is the same cohomology as the half-hex substitution. In fact, this tiling space is homeomorphic to the half-hex tiling space.

2.4. Rotations and other symmetries

A natural question about any pattern is "what are its symmetries?" An aperiodic tiling cannot have any translational symmetries, but it can have rotational or reflectional symmetries. We consider actions of reflection and rotation (and translation, of course) on the tiling space Ω, and examine how various quantities transform under that group action.

2.4.1. Decomposing by representation. Rotating a tile and then taking its boundary is the same as taking the boundary and then rotating. Likewise, rotations commute with *co*boundaries, and in most cases rotations commute with substitution, so it makes sense to decompose our cochain complexes, and the cohomology of our tiling space, into representations of whatever rotation group G acts on our tiling space. By Schurr's Lemma, neither the coboundary nor substitution can mix different representations, and our calculations can proceed one representation at a time.

The trouble with this approach is that representations are vector spaces, and our cochain complexes take values in $\mathbb{Z}$. We therefore consider the cohomology of tiling spaces with values in $\mathbb{R}$ rather than $\mathbb{Z}$. In the process we lose information about torsion and divisibility, but that's the price we have to pay.

For example, the tiles and substitution rules for the Penrose tilings are shown in Figures 4 and 5. There are four types of tiles, each in 10 orientations, and four types of edges, each in 10 orientations. There are only four kinds of vertices a, b, c, d each of which can sit in the center of a pattern with 5-fold rotational symmetry.

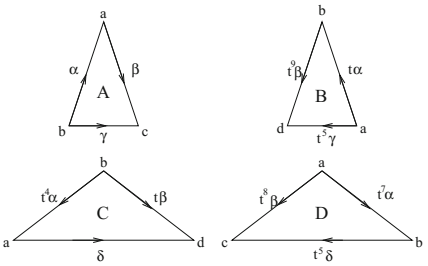

FIGURE 4. Four types of Penrose tiles

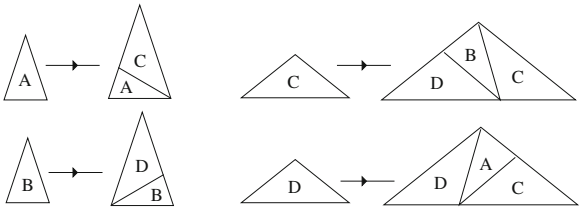

FIGURE 5. The Penrose substitution

This means that $a = t^2a$, $b = t^2b$, $c = t^2c$ and $d = t^2d$, where t is a rotation by $\pi/5$. In fact, $a = tb$ and $b = ta$, as can be seen from the fact that the α edge of A runs from b to a while the $t\alpha$ edge of B runs from a to b. Likewise, $c = td$ and $d = tc$. This tiling forces the border, so we do not need to collar.

The group $G = \mathbb{Z}_{10}$ acts on the Anderson–Putnam complex Γ by permuting the tiles, and the eigenvalues of the generator t are the 10th roots of unity. Each tile type, and each edge type, can be described by the module $\mathbb{R}[t]/(t^{10}-1)$. The polynomial $t^{10}-1$ factors as $(t-1)(t+1)(t^4+t^3+t^2+t+1)(t^4-t^3+t^2-t+1)$, with the factors corresponding to the primitive first, second, 5th and 10th roots, respectively. Each factor also corresponds to a representation. Since t^2 acts trivially on the vertices, only the representations with $t = \pm 1$ appear in C_0. More specifically, when working with the Anderson–Putnam complex, our chains complexes are:

$$\begin{aligned}
C_0(\Gamma) &= [\mathbb{R}[t]/(t^2-1)]^2 \\
&= [\mathbb{R}[t]/(t-1)]^2 \oplus [\mathbb{R}[t]/(t+1)]^2 \\
C_1(\Gamma) &= [\mathbb{R}[t]/(t^{10}-1)]^4 \\
&= \big[\mathbb{R}[t]/(t-1) \oplus \mathbb{R}[t]/(t+1) \oplus \mathbb{R}[t]/(t^4+t^3+t^2+t+1) \oplus \mathbb{R}[t] \\
&\qquad /(t^4-t^3+t^2-t+1)\big]^4 \\
C_2(\Gamma) &= [\mathbb{R}[t]/(t^{10}-1)]^4 \\
&= \big[\mathbb{R}[t]/(t-1) \oplus \mathbb{R}[t]/(t+1) \oplus \mathbb{R}[t]/(t^4+t^3+t^2+t+1) \oplus \mathbb{R}[t] \\
&\qquad /(t^4-t^3+t^2-t+1)\big]^4
\end{aligned} \tag{17}$$

The complexes $C^k(\Gamma)$ are the dual spaces of $C_k(\Gamma)$.

The matrices for the boundary maps $\partial_1 : C_1 \to C_0$ and $\partial_2 : C_2 \to C_1$ are:

$$\partial_1 = \begin{pmatrix} 1-t & -1 & -t & -1 \\ 0 & 1 & 1 & t \end{pmatrix}; \qquad \partial_2 = \begin{pmatrix} -1 & t & t^4 & -t^7 \\ -1 & t^9 & -t & t^8 \\ 1 & -t^5 & 0 & 0 \\ 0 & 0 & 1 & -t^5 \end{pmatrix} \tag{18}$$

in the representations $t = \pm 1$. In the other representations ∂_2 is the same, but ∂_1 is identically zero (since $C_0 = 0$). The coboundary maps δ_0 and δ_1 are the transposes of ∂_1 and ∂_2, only with t replaced by t^{-1}.

In the $t = 1$ representation, δ_1 has rank 2 and δ_0 has rank 1, and we get $H^0 = H^1 = \mathbb{R}$ and $H^2 = \mathbb{R}^2$. These are the elements of cohomology that are invariant under rotation. We say that this portion of the cohomology rotates like a scalar.

In the $t = -1$ representation, δ_0 and δ_1 are each rank 2, and we get $H^2 = \mathbb{R}^2$ and $H^1 = H^0 = 0$. This portion of the cohomology rotates like a *pseudoscalar*, flipping sign with every 36 degree rotation.

In the representation with $t^5 = 1$ but $t \neq 1$ (that is, with $t^4+t^3+t^2+t+1 = 0$), δ_1 is a rank-4 isomorphism, so all cohomologies vanish. In the representations with $t^5 = -1$ (but $t \neq -1$), δ_1 has rank 3, so $H^1 = H^2 = \mathbb{R}[t]/(t^4 - t^3 + t^2 - t + 1)$. This portion of the cohomology rotates like a *vector*, flipping sign after a 180 degree rotation.

Substitution acts on 2-cells by the matrix $\begin{pmatrix} t^7 & 0 & 0 & t^4 \\ 0 & t^3 & t^6 & 0 \\ t^3 & 0 & t^4 & 1 \\ 0 & t^7 & 1 & t^6 \end{pmatrix}$ and on 1-cells by $\begin{pmatrix} 0 & 0 & 0 & t^8 \\ t^4 & 0 & -t^7 & 0 \\ -t^7 & 0 & 0 & 0 \\ 0 & -t^3 & 0 & -t^3 \end{pmatrix}$. Both of these matrices are invertible for all representations (in fact, both have determinant 1), so $\check{H}^k(\Omega) = H^k(\Gamma)$ for $k = 0, 1, 2$. In summary:

$$\begin{aligned} \check{H}^0(\Omega) &= H^0(\Gamma) = \mathbb{R}[t]/(t-1) \\ \check{H}^1(\Omega) &= H^1(\Gamma) = \mathbb{R}[t]/(t-1) \oplus \mathbb{R}[t]/(t^4 - t^3 + t^2 - t + 1) \\ \check{H}^2(\Omega) &= H^2(\Gamma) = (\mathbb{R}[t]/(t-1))^2 \oplus (\mathbb{R}[t]/(t+1))^2 \oplus \mathbb{R}[t]/(t^4 - t^3 + t^2 - t + 1) \end{aligned} \tag{19}$$

The upshot is that $\check{H}^0(\Omega) = \mathbb{R}$ and is rotationally invariant, which is no surprise, since the generator is the constant function. $\check{H}^1(\Omega) = \mathbb{R}^5$, of which 4 dimensions rotate like vectors, with $t^5 = -1$, and one is rotationally invariant. $\check{H}^2(\Omega) = \mathbb{R}^8$, consisting of a rotationally invariant $\mathbb{R}^2$, a piece $\mathbb{R}^2$ that rotates like a pseudoscalar, and a piece $\mathbb{R}^4$ that rotates like a vector.

2.4.2. Three tiling spaces. For two-dimensional substitution like the Penrose tiling and the chair tiling, there are actually three tiling spaces to be considered. We have been considering the space Ω that is the (translational) orbit closure of a single tiling. This would be, for instance, the set of all chair tilings where the edges are parallel to the coordinate axes. We can also consider a larger space Ω_{rot} of all rotations of tilings in Ω. Finally we can consider the quotient space Ω_0 of tilings

modulo rotations. Ω_0 can either be viewed as $\Omega_{\rm rot}/S^1$ or as the quotient of Ω by the discrete group of rotations that acts on Ω. For the Penrose space, we would have $\Omega_0 = \Omega/\mathbb{Z}_{10}$, while for the chair tiling we would have $\Omega_0 = \Omega/\mathbb{Z}_4$.

The cohomologies of the three spaces are related as follows [ORS, BDHS]:

Theorem 2.7. *Working with real or complex coefficients, the cohomology of Ω_0 is isormorphic to the rotationally invariant part of the cohomology of Ω. The cohomology of $\Omega_{\rm rot}$ is isomorphic to the cohomology of $\Omega_0 \times S^1$.*

The upshot of this theorem is that Ω is the space with the richest cohomology, while $\Omega_{\rm rot}$ and Ω_0 have less cohomological structure. This is because all rotations on $\Omega_{\rm rot}$ are homotopic to the trivial rotation, and so act trivially on $\check{H}^*(\Omega_{\rm rot})$. Thus, only the rotationally invariant parts of the cohomology of Ω can manifest themselves in the cohomology of $\Omega_{\rm rot}$.

In general, $\Omega_{\rm rot}$ is not homeomorphic to $\Omega_0 \times S^1$, since the action of S^1 on $\Omega_{\rm rot}$ is typically not free. There are some tilings in $\Omega_{\rm rot}$ that have discrete k-fold rotational symmetry. For these tilings, rotation by $2\pi/k$ brings us back to the same tiling. $\Omega_{\rm rot}$ has the structure of a circle bundle over Ω_0 with some exceptional fibers corresponding to these symmetric tilings. (Seifert fibered 3-manifolds have a very similar structure.) When working with integer coefficients, these exceptional fibers can give rise to torsion in $\check{H}^2(\Omega_{\rm rot})$.[BDHS]

These relations can also be used to compute the cohomology of the pinwheel tiling space. When there are tiles that point in all directions, the only two well-defined spaces are $\Omega_{\rm rot}$ and Ω_0. For the pinwheel, the cohomology of Ω_0 can be computed with Barge–Diamond collaring, with the result that $\check{H}^0(\Omega_0) = \mathbb{Z}$, $\check{H}^1(\Omega_0) = \mathbb{Z}$ and $\check{H}^2(\Omega_0) = \mathbb{Z}[1/5] \oplus \mathbb{Z}[1/3] \oplus \mathbb{Z}^5 \oplus \mathbb{Z}_2$. This then determines the real cohomology of Ω_0, and, by Theorem 2.7, the real cohomology of $\Omega_{\rm rot}$. To compute the integer cohomology of $\Omega_{\rm rot}$, we have to consider the exceptional fibers. There are 6 pinwheel tilings with 2-fold rotational symmetry, as shown in Figure 9 below; these give rise to a $\mathbb{Z}_2^5$ term in a spectral sequence

$\mathbb{Z}$	$\mathbb{Z} \oplus \mathbb{Z}_2^5$	$\mathbb{Z}[1/5] \oplus \mathbb{Z}[1/3]^2 \oplus \mathbb{Z}^5 \oplus \mathbb{Z}_2$
$\mathbb{Z}$	$\mathbb{Z}$	$\mathbb{Z}[1/5] \oplus \mathbb{Z}[1/3]^2 \oplus \mathbb{Z}^5 \oplus \mathbb{Z}_2$

that computes the cohomology of $\Omega_{\rm rot}$. Furthermore, the d_2 map in the spectral sequence involves the torsion elements in a non-trivial way. The end result is that $\check{H}^1(\Omega_{\rm rot}) = \mathbb{Z}^2$, $\check{H}^2(\Omega_{\rm rot}) = \mathbb{Z}[1/5] \oplus \mathbb{Z}[1/3]^2 \oplus \mathbb{Z}^6 \oplus \mathbb{Z}_2^5$ and $\check{H}^3(\Omega_{\rm rot}) = \mathbb{Z}[1/5] \oplus \mathbb{Z}[1/3]^2 \oplus \mathbb{Z}^5 \oplus \mathbb{Z}_2$. For details of this calculation, see [BDHS] or [Sa3].

3. What is cohomology good for?

3.1. Distinguishing spaces

The most obvious use of topological invariants such as Čech cohomology is to distinguish spaces. If tiling spaces Ω and Ω' have cohomologies that are not isomorphic (as rings), then Ω and Ω' cannot be homeomorphic. If a group G of isometries of $\mathbb{R}^d$ (such as $\mathbb{Z}_{10}$ or $\mathbb{Z}_4$) acts on Ω and Ω', then we can decompose each cohomology group into representations of G. For each irreducible representation ρ of G, let $\check{H}^k_\rho(\Omega)$ be the part of $\check{H}^k(\Omega)$ that transforms under ρ. If Ω and Ω' are homeomorphic via a map that intertwines the action of G, then for each representation ρ we must have $\check{H}^k_\rho(\Omega) = \check{H}^k_\rho(\Omega')$. In particular, if a tiling space Ω' is related to the Penrose tiling Ω by a $\mathbb{Z}_{10}$-equivariant homeomorphism, then not only must $\check{H}^1(\Omega', \mathbb{R})$ equal $\mathbb{R}^5$, but $\check{H}^1(\Omega', \mathbb{R})$ must consist of a one-dimensional rotationally invariant piece and a four-dimensional piece that rotates like a vector.

For each subgroup $H < G$, we can also consider the topology of the set Ω_H of tilings in Ω that are fixed by H. If Ω and Ω' are tiling spaces with the same rotation group G, and if there exists an isomorphism that commutes with the action of G, then Ω_H and Ω'_H must be homeomorphic. In this sense, the structure of Ω_H is a topological invariant of the tiling space Ω. If $H_1 < H_2$, then $\Omega_{H_2} \subset \Omega_{H_1}$. The way that these different spaces nest within one another is also manifestly invariant.

If $d = 2$ and G is a subgroup of $SO(2)$, then Ω_H is not especially interesting. If N is the normalizer of H in G, then N acts on Ω_H, and there are typically only finitely many orbits. Understanding Ω_H boils down to counting these orbits and identifying how much symmetry a point in each orbit has.

Things get more interesting if $d > 2$, or if H involves reflections. In that case, there may be a subspace $V \subset \mathbb{R}^d$ whose vectors are fixed by the action of H. Ω_H is invariant under translation by elements of V, and can often be realized as a space of tilings of V, or as a disjoint union of several such lower-dimensional tiling spaces. In such cases, the Čech cohomology of Ω_H yields an interesting invariant.

We present two worked examples. We first consider the chair tiling of the plane, with $G = D_8 = O(2, \mathbb{Z})$, the group generated by rotation by 90 degrees and by reflection about the x axis. We compute the structure of Ω_H for every nontrivial subgroup $H < G$. We then consider the pinwheel tilings, with $G = O(2)$.

3.1.1. The chair tiling. We work with the "arrow" version of the chair tiling. This is a two-dimensional substitution tiling in which the tiles are all unit squares that meet full-edge to full-edge. Each square is decorated with an arrow pointing northeast, southeast, northwest or southwest. Rotation and reflection act naturally on arrows, so a counterclockwise rotation by 90 degrees would send a northeast arrow to a northwest, a northwest to a southwest, a southwest to a southeast, and a southeast to a northeast. Likewise, reflection about the x axis interchanges northeast and southeast arrows, and interchanges northwest and southwest arrows.

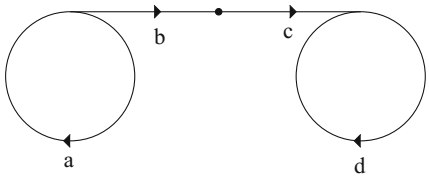

FIGURE 6. The approximant for Ω_{H_7}

The substitution on northeast arrows is

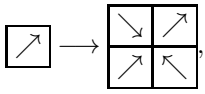

and the substitution on all other arrow tiles is obtained by rotating or reflecting this picture.

There are nine nontrivial subgroups of $G = D_8$. These include $H_1 = G$ itself, the rotation groups $H_2 = \mathbb{Z}_4$ and $H_3 = \mathbb{Z}_2$, the dihedral group $H_4 = D_4$ generated by reflections about the x and y axes, and the 2-element groups H_5 generated by reflection about the x axis, H_6 generated by reflection about the y axis, H_7 generated by reflection about the line $y = x$, and H_8 generated by reflection about the line $y = -x$. Finally, there is the dihedral group H_9 generated by H_7 and H_8.

There is only one tiling that is invariant under all of G, namely the fixed point of the substitution whose central patch involves four arrows pointing out from the origin: 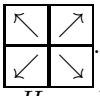 . This is also the only tiling that is invariant under H_2 or H_3 or H_4 or H_9. H_5 and H_6 are conjugate, so Ω_{H_5} and Ω_{H_6} are homeomorphic, with rotation by 90 degrees taking each set to the other. Likewise, Ω_{H_7} and Ω_{H_8} are homeomorphic. We therefore restrict our attention to Ω_{H_5} and Ω_{H_7}.

We begin with Ω_{H_5}. Since all vertices with incoming and outgoing arrows have either 3 or 0 incoming arrows, and since vertices with 3 incoming arrows cannot be symmetric under H_5, the tiles along the x axis must alternate between the form 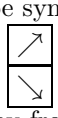 and 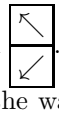 . Although the pattern along the x axis is periodic, there is a hierachy from the way that tiles along the x axis join with tiles once removed from the x axis to form clusters of four tiles, which join tiles even farther away to form clusters of 16, and so on. Ω_{H_5} is connected and is homeomorphic to the dyadic solenoid Sol_2, so $\check{H}^1(\Omega_{H_5}) = \mathbb{Z}[1/2]$.

To understand Ω_{H_7}, we look at symmetric configurations of tiles along the line $y = x$. All vertices take one of three forms:

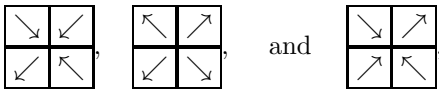

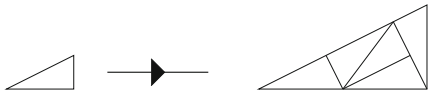

FIGURE 7. The pinwheel substitution

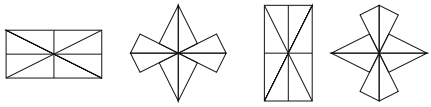

FIGURE 8. Central patches of pinwheel tilings with dihedral symmetry

and the second of these patterns can occur at most once. In other words, either all of the arrows along the line $x = y$ point northeast, or all point southwest, or all point outwards from a special point where four infinite-order supertiles meet.

Reading from southwest to northeast along the line $x = y$, there are four kinds of collared tiles that appear, which we label a, b, c and d. The label a means (SW)SW(SW), while b means (SW)SW(NE), c means (SW)NE(NE) and d means (NE)NE(NE). An a can be followed by an a or a b, a b is always followed by a c, a c is always followed by a c, and a d is always followed by a d. The Anderson–Putnam complex is then given by the "eyeglasses" graph shown in Figure 6.

Substitution sends edge a to aa, edge b to ab, edge c to cd and edge d to dd. The graph has $H^0 = \mathbb{Z}$ and $H^1 = \mathbb{Z}^2$, and substitution acts trivially on H^0 and by multiplication by 2 on H^1, so $\check{H}^0(\Omega_{H_7}) = \mathbb{Z}$ and $\check{H}^1(\Omega_{H_7}) = \mathbb{Z}[1/2]^2$.

One can apply a similar analysis to chair tilings in higher dimensions. For the three-dimensional chair tiling, the relevant group is the 24-element group G of symmetries of the cube, and there are significant subgroups of order 2, 3, 4, 6, and 12. Each cyclic subgroup H gives rise to a space Ω_H with non-trivial $\check{H}^1$, while the non-Abelian subgroups $H < G$ have Ω_H finite. For each non-Abelian H, the only invariant is $\check{H}^0(\Omega_H) = \mathbb{Z}^{|\Omega_H|}$.

3.1.2. The pinwheel tilings. The pinwheel tilings are based on a single tile, up to reflection, rotation and translation. It is a 1-2-$\sqrt{5}$ right triangle with substitution rule shown in Figure 7.

The maximal symmetry group for any pinwheel tiling is $H_1 = D_4$ (say, invariance under reflection about both the x and y axes). There are four such tilings, all closely related. Each is a fixed point of the square of the pinwheel substitution, with central patches shown in Figure 8.

There are two subgroups (up to conjugacy) of H_1, namely H_2 generated by 180 degree rotation, and H_3 generated by reflection about the x axis. There are six H_2-invariant tilings (plus rotations of the same), whose central patches are shown in Figure 9. All are periodic points of the substitution, the first four of period four and the last two of period two. In other words, Ω_{H_2} consists of six disjoint circles, so $\check{H}^1(\Omega_{\mathbb{Z}_2}) = \check{H}^0(\Omega_{\mathbb{Z}_2}) = \mathbb{Z}^6$. These six circles are the same exceptional fibers in the fibration $\Omega_{\text{rot}} \to \Omega_0$ that gave rise to torsion in $\check{H}^k(\Omega_{\text{rot}})$.

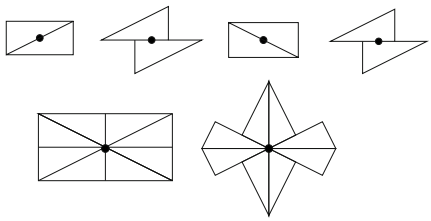

FIGURE 9. Central patches of pinwheel tilings with rotational symmetry

Finally, we consider tilings that are H_3-invariant. Since no tiles are themselves reflection-symmetric, there must be edges along the x axis, and these edges are either all hypotenuses or all of integer length.

For the symmetric tilings with hypotenuses along the x axis, we get a one-dimensional tiling space that comes from the substitution $a \to aabba$, $b \to baabb$, where a and b represent hypotenuses pointing in the two obvious directions. This edge substitution actually comes from the square of the pinwheel substitution, since the pinwheel substitution swaps hypotenuses and integer legs. A one-dimensional Barge–Diamond calculation shows that this set of tilings has $\check{H}^1 = \mathbb{Z}^2 \oplus \mathbb{Z}[1/5]$. Pinwheel substitution, applied only once, swaps this space with the space of symmetric tilings involving integer edges along the x axis, which therefore has the same cohomology. The upshot is that $H^1(\Omega_{H_3}) = \mathbb{Z}^4 \oplus \mathbb{Z}[1/5]^2$.

3.1.3. Asymptotic structures. A key difference between solenoids and spaces of one-dimensional non-periodic tilings is that tilings may be forward or backwards asymptotic. Suppose that $\mathbf{T}_1$ and $\mathbf{T}_2$ are tilings in the same tiling space Ω, but that the restrictions of $\mathbf{T}_1$ and $\mathbf{T}_2$ to the half-line $[0, \infty)$ are identical. Then $\mathbf{T}_1 - t$ and $\mathbf{T}_2 - t$ agree on a larger half-line $[-t, \infty)$, and $\lim_{t\to\infty} d(\mathbf{T}_1 - t, \mathbf{T}_2 - t) = 0$, where d is the metric on Ω. We say that $\mathbf{T}_1$ and $\mathbf{T}_2$ are *forward asymptotic*. Likewise, two tilings can be *backwards asymptotic*. The orbits of $\mathbf{T}_1$ and $\mathbf{T}_2$ are called *asymptotic composants*. Every substitution tiling space has a finite number of asymptotic composants, and the structure of these composants is reflected in the cohomology of Ω.

For instance, in the Thue–Morse tiling space there are four periodic points of the substitution of the form $\mathbf{T}_1 = \ldots a.a \ldots$, $\mathbf{T}_2 = \ldots a.b \ldots$, $\mathbf{T}_3 = \ldots b.a \ldots$ and $\mathbf{T}_4 = \ldots b.b \ldots$, where the central dot indicates the location of the origin. The tilings $\mathbf{T}_1$ and $\mathbf{T}_2$ are backwards asymptotic, as are $\mathbf{T}_3$ and $\mathbf{T}_4$. Likewise, $\mathbf{T}_1$ and $\mathbf{T}_3$ are forward asymptotic, as are $\mathbf{T}_2$ and $\mathbf{T}_4$. If we imagine asymptotic composants to be "joined at infinity", then the orbits of these four tilings form an *asymptotic cycle*. This asymptotic cycle manifests itself as the closed loop on Γ_{BD} that generated a $\mathbb{Z}$ term in $\check{H}^1(\Omega)$.

By studying asymptotic structures, Barge and Diamond [BD1] were able to construct a complete homeomorphism invariant of one-dimensional substitution tilings. Unfortunately, this invariant is extremely difficult to compute in practice.

As a practical alternative, Barge and Smith [BSm] constructed an *augmented cohomology* of one-dimensional substitution tilings. The precise definition involves the inverse limit of a variant of the Anderson–Putnam complex, but the basic idea is to identify all forward asymptotic tilings that are periodic points of the substitution, and separately to identify all backwards asymptotic periodic points. The cohomology of the resulting space, while not a complete invariant, yields finer information than the ordinary Čech cohomology.

In higher dimensions, asymptotic structures are more subtle, since there are (potentially) infinitely many directions to check. In 2 dimensions (with results that generalize somewhat to still higher dimensions), Barge and Olimb [BO] examined the *periodic branch locus* of a substitution, namely the set of pairs of tilings, each periodic under the substitution, that agree on at least a half-plane. From this locus, and from translates of these pairs along certain special directions, they construct a larger *branch locus* that can have a structure similar to that of a one-dimensional tiling space. The cohomology of the branch locus is a homeomorphism invariant of a tiling space.

With the chair tiling, as with a number of other examples, the branch locus seems to be closely related to the tilings that are symmetric under certain reflections, and the calculation of the cohomology of the branch locus resembles the computations of $\check{H}^1(\Omega_{H_5})$ and $\check{H}^1(\Omega_{H_7})$. These in turn are related to the quotient cohomology of the chair tiling space relative to the two-dimensional dyadic solenoid. While it might be a coincidence, all three computations seem to be telling the same story! Unfortunately, the general relation between cohomology of branch loci, cohomology of tilings with symmetry and quotient cohomology is not yet understood.

3.2. Gap labeling

For tilings of $\mathbb{R}^d$ with finite local complexity (with respect to translations), there is a natural trace map from the highest cohomology $\check{H}^d(\Omega)$ to $\mathbb{R}$. Each class $\alpha \in \check{H}^d(\Omega)$ can be represented by a pattern-equivariant d-cochain i_α. Pick any bounded region R of a tiling $\mathbf{T}$, let $i_\alpha(R)$ be the sum of the values of i_α on all of the tiles in R. If i_α and i'_α are cohomologous, then $i_\alpha - i'_\alpha = \delta i_\beta$ for some pattern-equivariant cochain i_β, and $i_\alpha(R) - i'_\alpha(R) = i_\beta(\partial R)$. Define

$$Tr(\alpha) = \lim_{r\to\infty} \frac{i_\alpha(B_r)}{\mathrm{Vol}(B_r)},$$

where B_r is the ball of radius r around the origin in a fixed tiling T. Since $\mathrm{Vol}(\partial B_r)/\mathrm{Vol}(B_r) \to 0$ as $r \to \infty$, different representatives for the class α yield the same limit. Likewise, if Ω is uniquely ergodic, then all tilings $\mathbf{T}$ yield the same limit.

For instance, in a Fibonacci tiling where the a tiles have length $\phi = (1+\sqrt{5})/2$ and the b tiles have length 1, there are on average ϕ a tiles for every b tile, so the indicator cochain i_a has trace $\phi/(\phi^2+1)$ and the cochain i_b has trace $1/(\phi^2+1)$. Since i_a and i_b generate $\check{H}^1$, the image of the trace map is $(\phi^2+1)^{-1}\mathbb{Z}[\phi]$.

The image of the trace map is called the *frequency module* of the tiling space. The frequency module is isomorphic to the *gap-labeling group*, which in K-theory is the image of a trace map in K^0. Besides being an invariant of topological conjugacies, the gap-labeling group is used (as the name implies) to label gaps in the spectra of Schrödinger operators associated with a tiling. The key theorem is due to Bellissard ([Bel], see also [BBG, BHZ, BKL]):

Theorem 3.1. *Let* $\mathbf{T}$ *be a tiling in a minimal and uniquely ergodic tiling space* X*, and let* $V : \mathbb{R}^d \to \mathbb{R}$ *be a strongly pattern-equivariant function. Consider the Schrödinger operator*

$$H = -\frac{\hbar^2}{2m}\Delta + V.$$

Let E_0 *be a point that is not in the spectrum of* H*. (That is,* E_0 *lies in a gap in the spectrum.) Then the integrated density of states up to energy* E_0 *is an element of the gap-labeling group of* Ω*.*

Elements of the frequency module (or gap-labeling group) should not be viewed as pure numbers. Rather, they have units of $(\text{Volume})^{-1}$, being the ratio of $i_\alpha(B_r)$ (a pure number) and $\text{Vol}(B_r)$. Likewise, the integrated density of states gives the number of eigenstates of H up to energy E_0 (a pure number) per unit volume.

Traces of cohomologies in all dimensions were studied in [KP], and are known as *Ruelle–Sullivan maps*. These traces give a ring homomorphism from $\check{H}^*(\Omega)$ to the exterior algebra of $\mathbb{R}^d$.

3.3. Tiling deformations

Some properties of a tiling are consequences of the geometry of the tiles, while others follow from the combinatorics of how tiles fit together. To distinguish between the two, we consider different tiling spaces that have the same combinatorics, and parametrize the possible tile shapes.

Let X be a tiling space. To specify the shapes of the tiles involved, we must indicate the displacement associated to every edge of every possible tile. Furthermore, if two tiles share an edge, then those two edges must be described by the same vector, and the vectors for all the edges around a tile must sum to zero.

In other words, the shapes of all the tiles is described by a co-closed vector-valued 1-cochain on a space obtained by taking one copy of each tile type and identifying edges that can meet. That is, a cochain on the Anderson–Putnam complex Γ_{AP}! Different geometric versions of the same combinatorial tiling space are described by different shape covectors on the same Anderson–Putnam complex.

Theorem 3.2 ([CS]). *Let* Ω *be a tiling space with shape cochain* α_0*. There is a neighborhood* U *of* α_0 *in* $C^1(\Gamma_{AP}, \mathbb{R}^d)$ *such that, for any two co-closed shape cochains* $\alpha_{1,2} \in U$*, the tiling spaces* Ω_{α_1} *and* Ω_{α_2} *obtained from* α_1 *and* α_2 *are mutually locally derivable (MLD) if and only if* α_1 *and* α_2 *are cohomologous.*

This theorem says that the first cohomology of Γ_{AP}, with values in $\mathbb{R}^d$, parametrizes local deformations of the shapes and sizes of the tiles, up to lo-

cal equivalence. By considering changes in the shapes and sizes of collared tiles and taking a limit of repeated collaring (as in Gähler's construction), we obtain

Theorem 3.3 ([CS]). *Infinitesimal shape deformations of tiling spaces, modulo local equivalence, are parametrized by the vector-valued cohomology* $\check{H}^1(\Omega, \mathbb{R}^d)$.

Among all shape changes, there are some that yield topological conjugacies. We call such a shape change a *shape conjugacy*. Shape conjugacies correspond to a subgroup $\check{H}^1_{an}(\Omega, \mathbb{R}^d)$ of $\check{H}^1(\Omega, \mathbb{R}^d)$ called the *asymptotically negligible* classes. These classes are neatly described in terms of pattern-equivariant functions:

Theorem 3.4 ([Kel3]). *A class in* $\check{H}^1(\Omega, \mathbb{R}^d)$ *is asymptotically negligible if and only if it can be represented as a strongly pattern-equivariant vector-valued* 1*-form that is the differential of a weakly pattern-equivariant vector-valued function.*

Asymptotically negligible classes don't just describe shape conjugacies. They essentially describe all topological conjugacies, thanks to

Theorem 3.5 ([KS1]). *If* $f : \Omega_X \to \Omega_Y$ *is a topological conjugacy of tiling spaces, then we can write* f *as the composition* $f_1 \circ f_2$ *of two maps, such that* f_1 *is a shape conjugacy and* f_2 *is an MLD equivalence.*

The importance of this theorem is that it allows us to check when a property of a tiling (e.g., having its vertices form a Meyer set) is invariant under topological conjugacies. One merely has to check whether the property is preserved by MLD maps (a local computation) and whether it is preserved by shape conjugacies. (The Meyer property turns out to be preserved by MLD maps but not by shape conjugacies[KS1].)

For substitution tilings, the asymptotically negligible classes are easy to characterize:

Theorem 3.6 ([CS]). *Let* Ω *be a substitution tiling space generated from a substitution* σ*. Let* σ^* *denote the action of* σ *on the vector space* $\check{H}^1(\Omega, \mathbb{R}^d)$*. The asymptotically negligible classes are the span of the generalized eigenvectors of* σ^* *with eigenvalues strictly inside the unit circle.*

For example, for the Fibonacci tiling we have $\check{H}^1(\Omega) = \mathbb{Z}^2$, so $\check{H}^1(\Omega, \mathbb{R}) = \mathbb{R}^2$. The substitution acts via the matrix $\left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, with eigenvalues $\lambda_1 = \phi$ and $\lambda_2 = 1 - \phi$ and eigenvectors $\left(\begin{smallmatrix} \lambda_{1,2} \\ 1 \end{smallmatrix}\right)$. Deformations proportional to the second eigenvector are asymptotically negligible, so all deformations are locally equivalent to an overall rescaling followed by an asymptotically negligible deformation. In particular, any two Fibonacci tiling spaces are topologically conjugate, up to an overall rescaling.

Similar arguments apply to any one-dimensional substitution tiling space where σ^* acts on $\check{H}^1$ via a Pisot matrix. The asymptotically negligible classes have codimension 1, so all deformations yield spaces that are topologically conjugate up to scale. In particular, for these substitutions, suspensions of subshifts (with all tiles having size 1) have the same qualitative properties as self-similar tilings.

For the Penrose tiling, $\check{H}^1(\Omega, \mathbb{R}^2) = \mathbb{R}^5 \otimes \mathbb{R}^2 = \mathbb{R}^{10}$. The eigenvalues of σ^* are ϕ and $1 - \phi$, each with multiplicity 4, and -1 with multiplicity 2. The multiplic-

ity 4 for the large eigenvalue corresponds to the four-dimensional family of linear transformations that can be applied to $\mathbb{R}^2$. (For self-similar tilings of $\mathbb{R}^d$, the leading eigenvalue of σ^* always has multiplicity d^2.) Meanwhile, the two deformations with eigenvalue -1 break the 180-degree rotational symmetry of the tiling space. Thus, any combinatorial Penrose tiling space that maintains 180-degree rotational symmetry must be topologically conjugate to a linear combination applied to the "standard" Penrose tiling space.

For cut-and-project tilings, the asymptotically negligible classes depend on the shape of the "window". When the window is not too complicated, there is an explicit description of these classes. This theorem applies even when $\check{H}^1(\Omega, \mathbb{R}^d)$ is infinite-dimensional.

Theorem 3.7 ([KS2]). *If the window of a cut-and-project scheme of codimension n is a polytope, or a finite union of polytopes, then* $\dim(\check{H}^1_{an}(\Omega, \mathbb{R}^d)) = nd$. *The elements of $\check{H}^1_{an}(\Omega, \mathbb{R}^d)$ correspond to projections from $\mathbb{R}^{n+d}$ to $\mathbb{R}^d$, and all shape conjugacies amount to simply changing the projection by which points in the acceptance strip are sent to $\mathbb{R}^d$.*

Besides the cohomology of strongly PE functions and forms, we can consider the *weak PE* cohomology of weakly PE functions and forms, and the *mixed* cohomology [Kel3]. We call a strongly PE form *weakly exact* if it can be written as d of a weakly PE $(k-1)$-form. The kth mixed cohomology $H^k_{PE,m}(\mathbf{T})$ of a tiling $\mathbf{T}$ is the quotient of the closed strongly PE k-forms by the weakly exact k-forms. This should not be viewed as a subgroup of $H^k_{PE}(\mathbf{T}) \equiv \check{H}^k(\Omega, \mathbb{R})$. Rather, it is a quotient of $H^k_{PE}(\mathbf{T})$ by those classes that can be representedy by weakly exact forms. This can be identified with a quotient of $\check{H}^k(\Omega, \mathbb{R})$. In dimension 1,

$$H^1_{PE,m}(\mathbf{T}) \equiv \check{H}^1(\Omega, \mathbb{R})/H^1_{an}(\Omega, \mathbb{R}).$$

Since $H^1_{an}(\Omega, \mathbb{R}^d)$ parametrizes shape conjugacies, this means that $H^1_{PE,m}(\mathbf{T})$ parametrized deformations of a tiling space $\Omega_{\mathbf{T}}$ *up to topological conjugacy* rather than up to MLD equivalence [Kel3].

3.4. Exact regularity

A measure on a tiling space is equivalent to specifying the frequencies of all possible patches. Specifically, let P be a patch in a specific location (say, centered at the origin). Let U be an open set in $\mathbb{R}^d$. Let $\Omega_{P,U}$ be the set of all tilings $\mathbf{T}$ such that, for some $x \in U$, $\mathbf{T} - x$ contains the patch P. In other words, $\mathbf{T}$ must contain the patch P at location x. As long as U is chosen small enough, there is at most one $x \in U$ that works. For any tiling $\mathbf{T}$, let $\mathrm{freq}_T(P)$ be the number of occurrences of P, per unit are, in $\mathbf{T}$. That is, restrict $\mathbf{T}$ to a large ball, divide by the volume of the ball, and take a limit as the radius goes to infinity. The ergodic theorem says that this limit exists for μ-almost every $\mathbf{T}$, with $\mathrm{freq}_{\mathbf{T}}(P) = \mu(\Omega_{P,U})/\operatorname{Vol}(U)$. If the tiling space is uniquely ergodic, then this statement applies to every $\mathbf{T}$, not just to almost every $\mathbf{T}$.

There are two natural questions. First, what are the possible values of $\text{freq}_{\mathbf{T}}(P)$? Second, as we consider larger and larger balls, how quickly does the number of occurrences of P per unit area approach $\text{freq}_{\mathbf{T}}(P)$? Both questions have cohomological answers.

Theorem 3.8. *For each patch P and each sufficiently small open subset U of $\mathbb{R}^d$, $\mu(\Omega_{P,U})/\operatorname{Vol}(U)$ takes values in the frequency module of X.*

Proof. Let i_P be a pattern-equivariant d-cochain that equals 1 on one of the tiles of P and is zero on all other tiles. Being of the top dimension, i_P is co-closed, and so represents a cohomology class. For any region R, $i_P(R)$ is just the number of occurrences of P in R. The limiting number per unit area $\text{freq}_{\mathbf{T}}(P)$ is then the trace of the class of i_P. □

Theorem 3.9 ([Sa4]). *Suppose that $\check{H}^d(\Omega, \mathbb{Q}) = \mathbb{Q}^k$ for some integer k. Then there exist patches $P_1, \ldots, P_k$ with the following property: for any other patch P, there exist rational numbers $c_1(P), \ldots, c_k(P)$ such that, for any region R in any tiling $\mathbf{T} \in X$, the number of appearances of P in R equals $\sum_{i=1}^{k} c_i(P) n_i(R) + e(P, R)$, where $n_i(R)$ is the number of appearances of P_i in R, and $e(P, R)$ is an error term computable from the patterns that appear on the boundary of R. In particular, the magnitude of $e(P, R)$ is bounded by a constant (that may depend on P) times the measure of the boundary of R. Furthermore, if $\check{H}^d(\Omega) = \mathbb{Z}^k$ is finitely generated over the integers, then we can pick the coefficients c_i to be integers.*

Corollary 3.10. *If the patches $P_1, \ldots, P_k$ have well-defined frequencies, then Ω is uniquely ergodic and there exist uniform bounds for the convergence of all patch frequencies to their ergodic averages. If the regions R are chosen to be balls, whose radii we denote r, then the number of P's per unit area approaches $\sum c_i \operatorname{freq}(P_i)$ at least as fast as one the frequency of one of the P_i's approaches $\operatorname{freq}(P_i)$, or as fast as r^{-1}, whichever is slower.*

Note that this theorem and its corollary apply to all tiling spaces, and not just to substitution tiling spaces.

Proof. Using the isomorphism between Čech and pattern-equivariant cohomology, pick patches $P_1, \ldots, P_k$ such that the cohomology classes of i_{P_i} are linearly independent. These classes then form a basis for $H^d_{PE}(X) = \mathbb{Q}^k$, and we can write $[i_P] = \sum c_i [i_{P_i}]$, where $[\alpha]$ denotes the cohomology class of the cochain α. This means that there is a $(d-1)$-cochain β such that $i_P = \sum_i c_i i_{P_i} + \delta\beta$. Then

$$\text{number of } P \text{ in } R = i_P(R) = \sum c_i i_{P_i}(R) + \delta\beta(R) = \sum c_i n_i(R) + \beta(\partial R). \quad (20)$$

Since β is pattern-equivariant there is a maximum value that it takes on any $(d-1)$-cell, so the error term $\beta(\partial R)$ is bounded by a constant times the area of the boundary of R. Dividing by the volume of R, the deviation of the left-hand

side from freq(P) is bounded by the deviation of $n_i(R)/\operatorname{Vol}(R)$ from freq(P_i) or by $|\partial R|/\operatorname{Vol}(R) \sim r^{-1}$.

If $\check{H}^d(\Omega) = \mathbb{Z}^k$, then the same argument applies with the patches chosen such that $[i_{P_i}]$ are generators of $\mathbb{Z}^k$ and with integral coefficients c_i. □

When Ω is a one-dimensional tiling space, it is possible to pick R such that ∂R is homologically trivial. Let β be pattern-equivariant with radius r_0, and let W be a word of length at least $2r_0$. Pick R to be an interval that starts in the middle of one occurrence of W and ends in the corresponding spot of another occurrence. Then $\delta\beta(R) = \beta(\partial R)$ vanishes, and the number of P's in R is *exactly* $\sum_i c_i(P)n_i(R)$. This is called *exact regularity* [BBJS, Sa4].

3.5. Invariant measures and homology

Exact regularity is dual to an earlier description [BG] of invariant measures in terms of the real-valued homology $H_d(\Gamma^n, \mathbb{R})$ of the approximants. Recall that measures do not pull back, but instead *push forward* like homology classes: Given a measure μ on a space X and a continuous map $f : X \to Y$, there is a measure $f_*\mu$ on Y. For any measurable set $S \subset Y$, $f_*\mu(S) := \mu(f^{-1}(S))$. Thus, a measure μ on a tiling space gives rise to a sequence of measures μ_n on the approximants Γ^n, with $(\rho_n)_*\mu_n = \mu_{n-1}$.

Integration gives a pairing between indicator d-cochains and measures. $\langle \mu, i_P \rangle = \text{freq}(P)$. This extends to a pairing between measures and cohomology, both for the tiling space and for each approximant. By the universal coefficients theorem, the dual space to the top cohomology group $H^d(\Gamma^n, \mathbb{R})$ is the top homology group $H_d(\Gamma^n, \mathbb{R})$, so we can view μ_n as living in $H_d(\Gamma^n, \mathbb{R})$, and μ as living in the inverse limit space $\varprojlim(H_d(\Gamma^n, \mathbb{R}), (\rho_n)_*)$. (The identification of μ_d as an element of $H_d(\Gamma^n, \mathbb{R})$ can also be seen more directly. A measure can be viewed as a chain satisfying certain "switching rules", or "Kirchoff-like laws". These rules are equivalent to saying that the boundary operator applied to μ_d is zero, i.e., that μ_d defines a homology class.)

The measure of any cylinder set is non-negative, so each μ_d must lie in the positive cone of $H_d(\Gamma^n, \mathbb{R})$. Not only is μ constrained to lie in the inverse limit of the top homologies of the approximants, but μ must lie in the inverse limit of the positive cones. All of the invariant measures on a tiling space can be determined from the transition matrices $(\rho_n)_*$, and in particular we can tell whether the tiling space is uniquely ergodic.

References

[AFHI] P. Arnoux, M. Furukado, E. Harriss and S. Itô, Algebraic numbers, free group automorphism and substitution on the plane, *Trans. Amer. Math. Soc.* **363** (2011) 4651–4699.

[AP] J.E. Anderson and I.F. Putnam, Topological invariants for substitution tilings and their associated C^*-algebras, *Ergodic Theory Dynam. Systems* **18** (1998), 509–537.

[Bel] J. Bellissard, The gap-labeling theorems for Schrödinger operators, in "From number theory to physics", pp. 538–630, Les Houches March 1989. J.M. Luck, P. Moussa and M. Waldschmidt eds.

[BBG] J. Bellissard, R. Benedetti, and J.-M. Gambaudo, Spaces of tilings, finite telescopic approximations and gap-labelling, *Comm. Math. Phys.* **261** (2006) 1–41.

[BBJS] M. Barge, H. Bruin, L. Jones and L. Sadun, Homological Pisot Substitutions and Exact Regularity, *Israel Journal of Mathematics* **188** (2012), 281–300.

[BD1] M. Barge and B. Diamond, A complete invariant for the topology of one-dimensional substitution tiling spaces, *Ergodic Theory Dynam. Systems* **21** (2001) 1333–1358.

[BD2] M. Barge and B. Diamond, Cohomology in one-dimensional substitution tiling spaces, *Proc. Amer. Math. Soc* **138** (2008) 2183–2191.

[BDHS] M. Barge, B. Diamond, J. Hunton and L. Sadun, Cohomology of Substitution Tiling Spaces, *Ergodic Theory and Dynamical Systems* **30** (2010) 1607–1627.

[BG] R. Benedetti, J.-M. Gambaudo, On the dynamics of G-solenoids: applications to Delone sets, *Ergodic Theory and Dynamical Systems* **23** (2003) 673–691.

[BHZ] J. Bellissard, D. Herrmann, and M. Zarrouati, Hulls of aperiodic solids and gap labeling theorems, in "Directions in mathematical quasicrystals", CRM monograph series **13** (2000) 207–259. M.B. Baake and R.V. Moody eds.

[BK] H. Boulezaoud and J. Kellendonk, Comparing different versions of tiling cohomology, *Topology and its applications* **157** (2010) 2225–2239.

[BKL] J. Bellissard, J. Kellendonk and A. Legrand, Gap labeling for three dimensional aperiodic solids, *C.R. Acad. Sci* **332**, Série I (2001) 521–525.

[BO] M. Barge and C. Olimb, Asymptotic structure in substitution tiling spaces, Ergodic Theory and Dynamical Systems **34** (2014) 55–94.

[BSa] M. Barge and L. Sadun, Quotient Cohomology for Tiling Spaces, *New York Journal of Mathematics* **17** (2011) 579–599.

[BSm] M. Barge and M. Smith, Augmented dimension groups and ordered cohomlogy, *Ergodic Theory and Dynamical Systems* **29** (2009) 1–35.

[BT] R. Bott and L. Tu, "Differential Forms in Algebraic Topology", Springer-Verlag, 1982.

[CS] A. Clark and L. Sadun, When Shape Matters: Deformations of Tiling Spaces, *Ergodic Theory and Dynamical Systems* **26** (2006) 69–86.

[Da] Inflation species of planar tilings which are not of locally finite complexity. *Proc. Steklov Inst. Math.* **230** (2002), 118–126.

[FHK] A. Forrest, J. Hunton and J. Kellendonk, Topological Invariants for Projection Method Patterns, Memoirs of the Amer. Math. Soc. **758**, (2002).

[F] N.P. Frank, A primer on substitution tilings of the Euclidean plane, *Expositiones Mathematicae* **26** (2008), 295–326.

[FR] N.P. Frank and E.A. Robinson, Jr., Generalized β-expansions, substitution tilings, and local finiteness, *Trans. Amer. Math. Soc.* **360** (2008) 1163–1177.

[FS] N.P. Frank and L. Sadun, Topology of (Some) Tiling Spaces without Finite Local Complexity, *Discrete and Continuous Dynamical Systems* **23** (2009) 847–865.

[Gah] F. Gähler, talk given at the conference "Aperiodic Order, Dynamical Systems, Operator Algebras and Topology", 2002, slides available at www.pims.math.ca/science/2002/adot/lectnotes/Gaehler

[GS] C. Goodman-Strauss, Matching Rules and Substitution Tilings, *Annals of Mathematics*, **147** (1998), 181–223.

[Hat] A. Hatcher "Algebraic Topology", (2002) Cambridge University Press.

[Kal] P. Kalugin, Cohomology of quasiperiodic patterns and matching rules, *J. Phys. A.* **38** (2005) 3115–3132.

[Kel1] Noncommutative geometry of tilings and gap-labelling, *Rev. Math. Phys.* **7** (1995), 1133–1180.

[Kel2] J. Kellendonk, Pattern-equivariant functions and cohomology, *J. Phys. A* **36** (2003) 5765–5772.

[Kel3] J. Kellendonk, Pattern-equivariant functions, deformations and equivalence of tiling spaces, *Ergodic Theory & Dynamical Systems* **28** (2008) 1153–1176.

[Ken] R. Kenyon, Self-Replicating Tilings, in "Symbolic dynamics and its applications", AMS Contemp. Math. Series 135, P. Walters ed., (1992) 239–263.

[KP] J. Kellendonk and I. Putnam, The Ruelle–Sullivan map for R^n actions, *Math. Ann.* **334** (2006), 693–711.

[KS1] J. Kellendonk and L. Sadun, Meyer sets, topological eigenvalues, and Cantor fiber bundles Journal of the London Math Society (2013), online 10.1112/jlms/jdt062

[KS2] J. Kellendonk and L. Sadun, Conjugacies of Model Sets, preprint 2014.

[Mos] B. Mossé: Puissances de mots et reconnaissabilité des point fixes d'une substitution. *Theor. Comput. Sci.* **99** (1992) 327–334.

[Moz] S. Mozes, Tilings, substitution systems and dynamical systems generated by them, *J. Analyse Math.* **53** (1989), 139–186.

[ORS] N. Ormes C. Radin and L. Sadun, A Homeomorphism Invariant for Substitution Tiling Spaces, *Geometriae Dedicata* **90** (2002), 153–182.

[Rad] C. Radin, The pinwheel tilings of the plane, *Annals of Math.* **139** (1994) 661–702.

[Sa1] L. Sadun, Tiling spaces are inverse limits, *J. Math. Phys.* **44** (2003) 5410–5414.

[Sa2] L. Sadun, Pattern Equivariant Cohomology with Integer Coefficients, *Ergodic Theory Dynan. Sys.* **27** (2007) 1991–1998.

[Sa3] L. Sadun, "Topology of Tiling Spaces", University Lecture Series, Vol. 46, American Mathematical Society, Providence, RI 2008.

[Sa4] L. Sadun, Exact Regularity and the Cohomology of Tiling Spaces, *Ergodic Theory and Dynamical Systems* **31** (2011) 1819–1834.

[SB] J. Savinien and J. Bellissard, A spectral sequence for the K-theory of tiling spaces, *Ergod. Th. & Dynam. Sys.* **29** (2009) 997–1031

[Sol] B. Solomyak, Nonperiodicity implies unique decomposition for self-similar translationally finite tilings. *Disc. & Comp. Geom.* **20** (1998), 265–279.

Lorenzo Sadun
Department of Mathematics
The University of Texas at Austin
Austin, TX 78712, USA
e-mail: sadun@math.utexas.edu

Progress in Mathematics, Vol. 309, 105–135
© 2015 Springer Basel

Spaces of Projection Method Patterns and their Cohomology

John Hunton

Abstract. We explain from the basics why the Čech cohomology of a tiling space can be realised in terms of group cohomology, and use this to explain how to compute the cohomology of a projection pattern.

Mathematics Subject Classification (2010). Primary: 52C23; Secondary: 37B50, 55C22, 54H20, 55R20.

Keywords. Aperiodic patterns, cut and project, model sets, cohomology, tilings, Cantor bundles.

1. Introduction

The aim of this article is provide a relatively simple explanation of the methods of Forrest, Gähler, Hunton and Kellendonk for the description and computation of the Čech cohomology of spaces of projection method tilings. The relevant published papers which present these methods and results include [5, 6, 7, 12, 10, 11] among others, but generally these either present computations with sparse detail of the theory behind them, or else present the theory in such full detail as to make the overarching ideas transparent to the reader only after significant study. This article attempts to provide a 'middle way', in which the important underlying ideas and structure of our approach is laid out, but many of the technical details are left for the interested reader to pursue in the original articles.

We take the opportunity also to cover a specific part of the programme that though addressed in part in [5], and at points in [7], has not had as complete treatment as it could have. This is the issue of why the Čech cohomology of a tiling space (any tiling space, but in particular that of a projection tiling) can be realised in terms of group homology or cohomology. This is a crucial part of the analysis of [7] *et al.* and one we intend to return to elsewhere.

The structure of the article is as follows. In Part I we address this last issue, explaining some of the underlying algebraic-topological approach that underpins

the main discussion of projection patterns, though this part is written in much greater generality. We explain why a space which can be described as a Cantor bundle over a torus has its Čech cohomology given by a certain group (co)homology. We know of course by [19] that in theory every tiling with translationally finite local complexity has this property, but to utilise this approach to actually compute requires the bundle structure to be given in a tangible and very describable manner: we show in Part II that for projection patterns this is the case[1].

In Part I we begin, Section 2, with a simple example, explaining how the cohomology of the space associated to a one-dimensional $\mathbb{Z}$ sub-shift can be described as a group cohomology, and in Section 3 consider the more general case of the cohomology of any space which can be presented as a Cantor bundle over a d-torus. The reader may also wish to read these sections in conjunction with the Appendix, which covers the basics of group homology and cohomology needed in this article. (This appendix also contains a short collection of the basic facts about Čech cohomology needed for the work in hand.)

In Part II we turn explicitly to projection patterns. Our discussion falls into three phases. In Section 4, after setting up the definition of what we mean by a *projection pattern*, we show how the defining data gives rise to a description of the associated pattern space Ω, and a certain map $\mu\colon \Omega \to \mathbb{T}$ to a high-dimensional torus. This is *not* the map giving the necessary bundle structure to Ω allowing us to use the machinery of Part I, but in Section 5 we use it to derive such a bundle structure, and give a couple of ways of describing $H^*(\Omega)$ in terms of a group homology $H_*(G; C)$ for certain groups G and coefficients C. As noted earlier, this in itself is not good enough in order to be able to compute unless one has a good description of the G module C and of the G action on it. This is the aim of Section 6 where we introduce a *filtration* of the fibre, described very explicitly in terms of the geometry of the initial projection scheme data. This section outlines the specific programme of computation for $H^*(\Omega)$ via this approach.

However, the reader may find it helpful to read Part II, especially this last section of it, alongside Part III Section 7 which presents in great detail the computation of an actual example (that of Ammann–Beenker). The results for this example have of course appeared several times in the literature already, but are presented here annotated with a full commentary, and will hopefully allow the reader to see the theory of Part II worked out very explicitly.

In addition to this explicit computation, Part III contains also Section 8 in which we collect a number of remarks and observations concerning the method described here, including some discussion of the complications which arise in more complex examples.

[1]There is an analogue here with the computation of the Čech cohomology of a tiling space in terms of an inverse limit. Again, we know from [8] and elsewhere that any tiling space can be written as an inverse limit of finite CW complexes $\Omega = \lim_{\leftarrow} X_r$, and hence its cohomology as the direct limit $H^*(\Omega) = \lim_{\rightarrow} H^*(X_r)$, but this does you little good unless you have a strong hold on the tower $\{X_r\}$, for example if the underlying tiling is generated by a substitution.

Perhaps it is worth noting what we do *not* discuss in this article. We do not discuss the many examples of projection patterns and their properties, applications or role in other aspects of the field. We do not discuss the more high-powered machinery needed to discuss the more complex computations, including the methods of machine computation that are frequently needed for all but the simplest cases, or the more sophisticated homological algebra that can be utilised: for these the reader should consult respectively [11] §5 and §3. We only hint (Section 6) at Kalugin's approach [15] to the subject and how this fits into the scheme described here, again a far more detailed discussion of this point can be found in [11] §3 and §4. Finally, we do not look at the non-commutative geometry approach to these patterns, though the reader interested in computing the K-theory of these patterns will find much in the current work to aid efforts, at least for patterns in small (< 4) dimensions. (Again, see [11], end of §6, and Putnam's work [18].)

Part I: Cantor dynamics and cohomology

2. A very simple example

The purpose of this part is to sketch the perspective from which the cohomology of a tiling space may be seen as a group cohomology. Our particular interest, where this becomes an especially useful viewpoint, is with the projection tilings of Part II, but here we illustrate the general principle, starting in this section with an elementary, one-dimensional example.

Suppose we have a one-dimensional $\mathbb{Z}$ sub-shift, say generated by a bi-infinite word T in some finite alphabet $\mathcal{A}$. For now, we make no assumptions about properties such as repetitively, or even aperiodicity. Let us think of T as a tiling of $\mathbb{R}$ by unit length tiles, each labeled by an element of $\mathcal{A}$. In the usual fashion, as detailed elsewhere in this volume, we can construct a 'tiling space', Ω_T, for example as the completion under one of the usual metrics of the set $T + \mathbb{R}$ of all translated images of T. Let us denote by X a copy of, say, the transversal given by those tilings in Ω_T with tile boundaries at integer points, a totally disconnected space. Translation of a tiling in Ω_T by an integer defines a $\mathbb{Z}$-action on Ω_T:

$$\psi\colon \mathbb{Z} \times \Omega_T \to \Omega_T$$

which restricts to a $\mathbb{Z}$ action on X, which we also denote by ψ.

There is an obvious map $\pi\colon \Omega_T \to \mathbb{R}/\mathbb{Z} = \mathbb{S}^1$ to the circle given by declaring $\pi(t)$ to be the position of the origin within (any) tile in the tiling $t \in \Omega_T$. This gives Ω_T a structure of a fibre bundle over the circle, with fibre X. Clearly π commutes with the action of $\mathbb{Z}$ as given by ψ on Ω_T and the trivial action on $\mathbb{S}^1$; the action ψ on X is the monodromy action of this bundle (i.e., the action of the fundamental group of the base space $\mathbb{S}^1$ on the total space).

Now consider $\mathbb{S}^1$ cut into two (closed) semicircles, and denote by A and B their pre-images in Ω_T under π. The total disconnectedness of X means that each

of A and B are just copies of $I \times X$ for a closed interval I, and so, up to homotopy equivalence, they are just copies of X. Their intersection is a disjoint pair of copies of X, and the resulting Mayer–Vietoris decomposition of Ω_T gives a long exact sequence in Čech cohomology

$$\cdots \to H^n(\Omega_T) \to H^n(X) \oplus H^n(X) \xrightarrow{\Psi} H^n(X) \oplus H^n(X) \to H^{n+1}(\Omega_T) \to \cdots . \tag{2.1}$$

The map Ψ is given by the matrix $\begin{pmatrix} 1 & 1 \\ 1 & \psi^* \end{pmatrix}$, where ψ^* denotes the homomorphism in $H^n(-)$ given by the map

$$X \to X, \qquad x \mapsto \psi(1, x) .$$

Of course the total disconnectedness of X means that $H^n(X)$ is only non-trivial if $n = 0$ and there are only 4 potentially non-trivial terms in our sequence (2.1).

The nature of the matrix Ψ means that we can simplify this sequence: for example,

$$\ker \Psi = \{(a, b) \in H^n(X) \oplus H^n(X) \,|\, (a + b, a + \psi^*(b)) = (0, 0)\}$$

that is, $\{(a, b) \,|\, a = -b,\ a - \psi^*(a) = 0\}$ and so we can identify

$$\ker \Psi \cong \ker (1 - \psi^* \colon H^n(X) \to H^n(X)) .$$

Similarly, $\operatorname{coker} \Psi$ may be identified with $\operatorname{coker} (1 - \psi^*) \colon H^n(X) \to H^n(X))$ and (2.1) reduces to a long exact sequence

$$0 \to H^0(\Omega_T) \longrightarrow H^0(X) \xrightarrow{(1-\psi^*)} H^0(X) \longrightarrow H^1(\Omega_T) \to 0 ,$$

a structure not unrelated to the Pimsner-Voiculescu sequence in K-theory for crossed products of algebras by $\mathbb{Z}$. In particular, we obtain

Proposition 2.1.

$$\begin{aligned} H^0(\Omega_T) &= \ker(1 - \psi^*) \\ &= \textit{the } \psi^*\textit{-invariant elements of } H^0(X) \\ &= H^0(\mathbb{Z}; H^0(X)) \end{aligned}$$

$$\begin{aligned} H^1(\Omega_T) &= \operatorname{coker} (1 - \psi^*) \\ &= \textit{the } \psi^*\textit{-coinvariant elements of } H^0(X) \\ &= H^1(\mathbb{Z}; H^0(X)) \end{aligned}$$

The identifications in this statement of the group cohomologies $H^0(\mathbb{Z}; -)$ and $H^1(\mathbb{Z}; -)$ as invariants and coinvariants of the arguments are elaborated on in the Appendix. We note there also that the Čech cohomology $H^0(-)$ may be interpreted as the set $C(-; \mathbb{Z})$, of continuous functions from the argument space to $\mathbb{Z}$.

In the case of T aperiodic and repetitive, the transversal X is a Cantor set and the action ψ is minimal. In this situation, the only ψ-invariant continuous

functions $X \to Z$ are the constant ones, and we recover the well-known fact that $H^0(\Omega_T) = \mathbb{Z}$. The other cohomology group, $H^1(\Omega_T)$, identified as the group

$$C(X;\mathbb{Z})/_{\langle f - f\circ\psi\rangle}\,,$$

may be a good deal more complicated.

In the case that T is periodic, say of period m, the transversal X consists of m points, which if we identify with the points of $\mathbb{Z}/m$, realises the action of ψ as the usual (quotient) action of $\mathbb{Z}$ on $\mathbb{Z}/m$. In this case, it is easy to check that the only ψ-invariant functions $\mathbb{Z}/m \to \mathbb{Z}$ are the constant ones, and again $H^0(\Omega_T) = \mathbb{Z}$. For the coinvariants, note that $C(\mathbb{Z}/m;\mathbb{Z})$ is generated by the m indicator functions χ_p, $p \in \mathbb{Z}/m$, that is, χ_p denoting the function $\mathbb{Z}/m \to \mathbb{Z}$ taking value 1 on $p \in \mathbb{Z}/m$ and 0 elsewhere. The action ψ gives $\chi_{p+1} = \chi_p \circ \psi$, so in the group of coinvariants, all the χ_p are equivalent, and we obtain $H^1(\Omega_T) = \mathbb{Z}$, which of course should be expected since Ω_T is circle and the map π is the m-fold cover of $\mathbb{S}^1$.

3. $\mathbb{Z}^d$ dynamics

The data involved in the model of the tiling space for the one-dimensional subshift in the last section amounted to a disconnected transversal X with a homeomorphism ψ. Such data is equivalent to an action of the group $\mathbb{Z}$ on X via the map

$$\mathbb{Z} \times X \longrightarrow X \qquad (n,x) \mapsto \psi^{(n)}(x)\,.$$

In turn, this information is equivalent to a fibre bundle

$$\mathbb{M}_\psi \longrightarrow \mathbb{S}^1$$

where $\mathbb{M}_\psi$ denotes the *mapping torus* of $\psi\colon X \to X$, the quotient space

$$\mathbb{M}_\psi = I \times X/_{(0,x)\sim(1,\psi(x))} = \mathbb{R} \times X/_{(r,x)\sim(r+1,\psi(x))}\,.$$

In the previous section, $\mathbb{M}_\psi$ was what we were calling Ω_T.

In this section we turn attention to computations for higher-dimensional actions. Bearing in mind these three equivalent ways of encoding the data we need to consider, we suppose given a system, perhaps a tiling space Ω, that has the structure of a fibre bundle over a d-torus $\mathbb{T}^d$, again with totally disconnected[2] fibre (transversal) X. Again, this is equivalent to having an action of $\mathbb{Z}^d$ on X

[2]The reader may reasonably wonder about the assumption that the transversal X be totally disconnected, especially if, for example, she is interested in tilings with infinite local complexity (ILC) where the transversal could indeed have a richer topological structure. In short, yes, the ideas here can indeed be adapted to cover this case, but as might be expected, at the expense of becoming considerably more complicated. The first complication (assuming of course that one already has one's tiling space described as a fibre bundle over $\mathbb{T}^d$ in any sort of useful way) is when the Čech cohomology $H^*(X)$ of the transversal is non-trivial in positive dimensions, and this extra cohomology clearly needs to be taken into account, but then later exact sequences in what follows here will need to be replaced by spectral sequences, with potentially non-trivial differentials and extensions. We intend to address these issues further elsewhere.

(the monodromy action), and this is equivalent to having d commuting homeomorphisms $\psi = \{\psi_1, \ldots, \psi_d\}$ acting on X.

Let us denote again by $\mathbb{M}_\psi$ the mapping torus in this situation: we define $\mathbb{M}_\psi$ as the quotient space $\mathbb{R}^d \times X/\sim$ for the equivalence relation

$$((r_1, \ldots, r_i, \ldots, r_d), x) \sim ((r_1, \ldots, r_i + 1, \ldots, r_d), \psi_i(x)), \qquad i = 1, \ldots, d.$$

(Alternatively, as the corresponding quotient of $I^d \times X$.)

The fibre bundle structure $\mathbb{M}_\psi \xrightarrow{\pi} \mathbb{T}^d$ is then induced by the projection mapping the transversal X to a single point

$$\mathbb{R}^d \times X/\sim \to \mathbb{T}^d \qquad [(r_1, \ldots, r_i, \ldots, r_d), x] \mapsto [r_1, \ldots, r_i, \ldots, r_d].$$

The fibre (pre-image of a point in $\mathbb{T}^d$) is then of course X.

Theorem 3.1 ([5]). *There is an isomorphism $H^n(\mathbb{M}_\psi) \cong H^n(\mathbb{Z}^d; C(X; \mathbb{Z}))$, between the Čech cohomology of the mapping torus and the group cohomology on the right-hand side.*

Sketch of proof. The 'high-tech' proof of this is as a straightforward application of the Serre spectral sequence (SSS) for the fibre bundle $X \to \mathbb{M}_\psi \to \mathbb{T}^d$, and using the fact, as in the previous section, that $H^n(X) = C(X; \mathbb{Z})$ for $n = 0$ and vanishes for positive n (a point where we use the total disconnectedness of X).

However, a more 'bare hands' explanation is also possible (which actually is really just the SSS again, written out from first principles). First consider $\mathbb{T}^d$ as a cell complex, say with the usual product structure of 1 0-cell, d 1-cells, ..., $\binom{d}{n}$ n-cells, and so on. The torus is the union of these cells. Let us write $\mathbf{R}_*$ for the associated chain complex. Then the (dual) cochain complex

$$0 \to R^0 \to R^1 \to \cdots \to R^d \to 0$$

where R^n denotes $H^*(\text{the } n\text{-cells}) = \hom(R_n; \mathbb{Z})$, the cohomology (all in degree 0) of the disjoint union of the n-cells (of course, $R^n = \bigoplus_{n\text{-cells}} \mathbb{Z}$) encodes the construction of $\mathbb{T}^d$ from these cells, with the 'gluing' information on how to put the cells together coded in the coboundary maps $R^n \to R^{n+1}$. Further, this cochain complex computes, by taking its homology, the cohomology $H^*(\mathbb{T}^d)$ (for any cohomology theory, Čech or otherwise, as this is a finite cell complex, so all theories are equivalent).

Now consider the induced decomposition of the space $\mathbb{M}_\psi$, where we think of the component pieces being the π-preimage of each n-cell in $\mathbb{T}^d$. Let us write A_n for the disjoint union of the pieces $\pi^{-1}(\lambda)$ for each n-cell λ; as for the torus, $\mathbb{M}_\psi$ can be constructed by gluing together the A_n's with appropriate data, given by the same data as that needed for gluing the boundaries of the n-cells in $\mathbb{T}^d$, together with the monodromy information provided by the action of the ψ_i's.

As each $\pi^{-1}(\lambda)$ is, up to homotopy equivalence, just a copy of the fibre X (the n-cell λ being contractable), we see that

$$H^*(A_n) = \bigoplus_{n\text{-cells}} H^*(X) = \bigoplus_{n\text{-cells}} C(X; \mathbb{Z}) \qquad \text{concentrated in degree 0.}$$

Putting this information together, we have the cohomology of $\mathbb{M}_\psi$ given as the homology of the diagram

$$0 \to \bigoplus_{\text{0-cells}} C(X;\mathbb{Z}) \to \bigoplus_{\text{1-cells}} C(X;\mathbb{Z}) \to \cdots \to \bigoplus_{d\text{-cells}} C(X;\mathbb{Z}) \to 0$$

which is just the homology of $\hom_{\mathbb{Z}\mathbb{Z}^d}(\mathbf{R}_* \otimes_{\mathbb{Z}} \mathbb{Z}\mathbb{Z}^d; C(X;\mathbb{Z}))$, that is, by definition, the group cohomology $H^*(\mathbb{Z}^d; C(X;\mathbb{Z}))$ (see Appendix), by virtue of the fact that our exact sequence $\mathbf{R}_* \otimes_{\mathbb{Z}} \mathbb{Z}\mathbb{Z}^d$ is a projective (in fact free) resolution for $\mathbb{Z}^d$. In fact it is precisely the resolution discussed in Example 9.9. □

As in the previous section, the identification of the cohomology $H^*(\mathbb{M}_\psi)$ with this group cohomology allows some immediate corollaries, just from the nature of group cohomology of free abelian groups (again, see the Appendix).

Corollary 3.2.

1. $H^n(\mathbb{M}_\psi)$ *is non-trivial only for* $0 \leqslant n \leqslant d$.
2. $H^0(\mathbb{M}_\psi)$ *is the group of* $\mathbb{Z}^d$*-invariant elements of* $C(X;\mathbb{Z})$. *If the* $\mathbb{Z}^d$*-action is minimal, this is just a single copy of* $\mathbb{Z}$.
3. $H^d(\mathbb{M}_\psi)$ *is the group of* $\mathbb{Z}^d$*-coinvariant elements of* $C(X;\mathbb{Z})$.

Remark 3.3. If ψ' denotes some subset of $\psi = \{\psi_1, \ldots, \psi_d\}$, say a collection of $0 \leqslant n < d$ of the homeomorphisms, then there is also a fibre bundle structure

$$\mathbb{M}_\psi \xrightarrow{\pi'} \mathbb{T}^{d-n}$$

with fibre $\mathbb{M}_{\psi'}$. We demonstrate this in the special case where $n = d-1$, but the general case is very similar. So, suppose ψ' leaves out precisely one of the ψ_i, say the last one, $\psi' = \{\psi_1, \ldots, \psi_{d-1}\}$. Then the fibration

$$\mathbb{M}_{\psi'} \longrightarrow \mathbb{M}_\psi \xrightarrow{\pi'} \mathbb{S}^1 \tag{3.1}$$

can be viewed as an iterated construction of $\mathbb{M}_\psi$ given by first constructing the mapping torus for the first $d-1$ homeomorphisms, to obtain $\mathbb{M}_{\psi'}$, then forming the full mapping torus as

$$\mathbb{M}_\psi = \mathbb{R} \times \mathbb{M}_{\psi'}/\!\sim' \qquad \begin{array}{l}\text{where } \sim' \text{ is given by} \\ (r, [(r_1, \ldots, r_{d-1}), x]) \sim' (r+1, [(r_1, \ldots, r_{d-1}), \psi_d(x)])\end{array}$$

and the map π' given by sending $\mathbb{M}_{\psi'}$ to a point.

This provides an iterative approach to the computation of the groups $H^*(\mathbb{M}_\psi)$. Supposing one has a computation for the cohomology of the mapping torus associated to the first $d-1$ homeomorphisms, the bundle (3.1), via a decomposition as in Section 2, gives short exact sequences

$$\begin{aligned} 0 \to \operatorname{coker}\left[H^{n-1}(\mathbb{M}_{\psi'}) \xrightarrow{(1-\psi_d^*)} H^{n-1}(\mathbb{M}_{\psi'})\right] \longrightarrow H^n(\mathbb{M}_\psi) \longrightarrow \\ \longrightarrow \ker\left[H^n(\mathbb{M}_{\psi'}) \xrightarrow{(1-\psi_d^*)} H^n(\mathbb{M}_{\psi'})\right] \to 0\,. \end{aligned} \tag{3.2}$$

If we write $\mathcal{C}_i$ and $\mathcal{I}_i$ for respectively the coinvariant and invariant functors with respect to the action ψ_i^* on the relevant modules, the sequence (3.2) becomes

$$0 \to \mathcal{C}_d(H^{n-1}(\mathbb{M}_{\psi'})) \to H^n(\mathbb{M}_\psi) \to \mathcal{I}_d(H^n(\mathbb{M}_{\psi'})) \to 0$$

and this iterates to express $H^n(\mathbb{M}_\psi)$ via a series of extensions of the form

$$\mathcal{F}_d\mathcal{F}_{d-1}\cdots\mathcal{F}_1(C(X;\mathbb{Z}))$$

where each $\mathcal{F}_i$ is either $\mathcal{C}_i$ or $\mathcal{I}_i$, and exactly n of them are $\mathcal{C}$'s. The reader should be warned however that the extensions in this iteration, such as (3.2) do not necessarily split. However, if one is working with coefficients over a field, such as with rational cohomology, then all the extensions will be trivial, and we obtain

Theorem 3.4.

$$H^n(\mathbb{M}_\psi;\mathbb{Q}) \cong \bigoplus_{\textit{exactly } n\ \mathcal{C}\textit{'s}} \mathcal{F}_d\mathcal{F}_{d-1}\cdots\mathcal{F}_1(C(X;\mathbb{Z})\otimes\mathbb{Q})\,.$$

Remark 3.5. Before leaving this part, let us reiterate that the group cohomological interpretation for the cohomology of a tiling space Ω we have sketched has relied on first representing Ω as a fibre bundle over a torus, presenting the cohomology in terms of the cohomology of the fibre and the monodromy action of $\mathbb{Z}^d$, the fundamental group of the base space. As a practical method for computing cohomology, this then stands or falls on being able to have a practical description of Ω as such a bundle and of the corresponding action ψ. While we know theoretically [19] that a wide class of tilings can be written (at least up to homeomorphism, which is quite good enough for cohomological calculations) as fibre bundles over tori, the need for a tangible description of the monodromy, good enough to be able to use to do calculations, is a much bigger demand. The essence of Part II of this article is that, remarkably, for cut and project patterns one *does* have good descriptions of this information, and computations are certainly possible.

Part II: Cut and project patterns and their pattern spaces

4. Projection method patterns

We now turn to consideration of cut and project patterns, their pattern spaces and the analysis of their cohomology. In this section we begin with a description of the basic paraphernalia needed for a cut and project pattern, explain how this data gives a direct construction of a large part of the pattern space Ω, and passing to a completion, to the whole of Ω.

The following data is needed to define such a pattern.

Definition 4.1. A *cut and project scheme* consists of data

1. $\mathcal{E}$ an N-dimensional Euclidean space;

2. $\Gamma \subset \mathcal{E}$ a discrete, cocompact subgroup of $\mathcal{E}$, which we shall also refer to as our *lattice*. The quotient space $\mathcal{E}/\Gamma$ is then an N-dimensional torus, which we denote $\mathbb{T}$.
3. A linear decomposition $\mathcal{E} = E \oplus F$ into subspaces $E \cong \mathbb{R}^d$ and $F \cong \mathbb{R}^n$ (so $N = d + n$), which we assume to be *irrationally positioned* with respect to Γ, meaning that the linear projections $\pi_E \colon \mathcal{E} \to E$ and $\pi_F \colon \mathcal{E} \to F$ are each one-to-one and with dense image on the lattice Γ.
4. An *acceptance domain* $K \subset F$ which, for our purposes here, will always be a finite union of compact, non-degenerate polyhedra[3] in F. The boundary ∂K of K consists of a union of finite $(n-1)$-dimensional faces; we use f_i to denote such a face, with i running over some finite indexing set.

In what follows it will be convenient to define also the associated *acceptance strip*, the subspace $\Sigma = K + E \subset \mathcal{E}$. We shall say that the strip Σ, or the acceptance domain K, is in *non-singular position* if $\Gamma \cap (\partial K + E) = \emptyset$, i.e., the boundary of the strip contains no lattice points, or, equivalently, if $\pi_F(\Gamma) \cap \partial K = \emptyset$.

Definition 4.2. The *cut and project point pattern* defined by the data above is the subset

$$\mathcal{P} = \{\pi_E(x) \,|\, x \in \Sigma \cap \Gamma\} \subset E\,.$$

Alternatively, but equivalently, $\mathcal{P}$ can be considered as the set $E \cap (\Gamma - K)$. Clearly $\mathcal{P}$ is a subset of the d-dimensional space E; we refer to d as the *dimension* of the pattern, and n, the dimension of F, as its *codimension*.

These notions were used respectively, for example, in the articles of Forrest, Hunton and Kellendonk [7] and of Kalugin [15] in their analyses of cohomology, but other variations exist too, such as the dual method using Laguerre complexes which is more elegant for some tilings such as the Penrose tilings, and was utilised in [6]. See Moody's work, for example [17], for a wide ranging discussion of approaches to cut and project patterns and their properties.

Remark 4.3. The data and constructions given automatically imply that the patterns generated are aperiodic and of finite local complexity. In particular, aperiodicity follows directly from the irrational position assumption. The patterns formed are however not necessarily repetitive, and it can be checked that they are repetitive if and only if the acceptance strip Σ is in non-singular position.

Remark 4.4. Note that the lattice Γ clearly acts on $\mathcal{E}$ by translation; it also acts, via the projections π_E, π_F on E and F. The latter will play an important role in this article, and we note explicitly the action as given by

$$\gamma \cdot x := x + \pi_F(\gamma) \qquad \text{for } \gamma \in \Gamma \text{ and } x \in F.$$

[3] In this article, a *polyhedron* in an n-dimensional euclidean space R is a compact subspace with non-empty interior given as the intersection of a finite number of *half-spaces*, where a half-space means a subset of R consisting of all points on or to one side of some affine $(n-1)$-dimensional linear subspace of R.

Unless necessary, we shall not usually distinguish notationally between the Γ action the whole space $\mathcal{E}$ and the (projected) Γ action on F.

In fact, a cut and project scheme, that is, the data given by $\Gamma \subset \mathcal{E} = E \oplus F$ and the *translation class* in $\mathcal{E}$ of the acceptance strip Σ defines a whole parameterised family of point patterns in E: for each choice of where we embed Σ in $\mathcal{E}$ we obtain a, possibly distinct, point pattern. One way of describing this is as follows.

Definition 4.5. Suppose given the data $\Gamma \subset \mathcal{E} = E \oplus F$ and a particular acceptance domain $K \subset F$, as in Definition 4.1. For each point $v \in \mathcal{E}$ define the point set

$$\begin{aligned} P_v &= \{\pi_E(x) \,|\, x \in \Sigma \cap (\Gamma + v)\} \\ &= E \cap (\Gamma + v - K). \end{aligned}$$

Remark 4.6. It is worth noting that the usual translation action on a point pattern has a simple representation in this perspective. Suppose $P_v \subset E$ is some point pattern given by this data, and $x \in E$. Then the x-translate of the pattern P_v is just P_{v+x}.

Remark 4.7. We note also that the pattern P_v depends only on the class of v in $\mathcal{E}/\Gamma = \mathbb{T}$. In fact $P_v = P_{v'}$ if and only if $v - v' \in \Gamma$. In particular, it will be useful to define $q\colon \mathcal{E} \to \mathbb{T}$ as the quotient map

$$q\colon \mathcal{E} \to \mathcal{E}/\Gamma\,.$$

Thus the points of $\mathbb{T}$ paramaterise a large class of patterns, but these are not exactly the set of patterns comprising the pattern space Ω of any given P_v: in fact not all of these patterns in this parameterisation will necessarily be points in Ω, and nor does every point in Ω appear in this parameterisation. Nevertheless, we will use this parameterisation to conveniently model the pattern space Ω; key will be a distinction between two classes of types of points $v \in \mathcal{E}$.

Definition 4.8. Define the set S of *singular points* in $\mathcal{E}$ as

$$S = \{v \in \mathcal{E} \,|\, \pi_F(v) \in \partial K + \pi_F(\Gamma)\} = E + \Gamma + \partial K\,.$$

Denote by NS its complement, the set of *non-singular points*. In the terminology above, the acceptance strip Σ is in non-singular position if and only if $0 \in NS$, in fact, if and only if $E \subset NS$.

Remark 4.9. It may be checked that if $v \in NS$ then the corresponding P_v is repetitive; moreover, two patterns P_v and $P_{v'}$ with both v and $v' \in NS$ are locally indistinguishable, in the sense that arbitrarily large, compact patches in one occur also in the other, and hence each pattern lies in the other's pattern space. These statements are generally not true without the requirement that the parameters v, v' are non-singular.

The upshot of this is that if $v \in NS$, with its coset $v+\Gamma$ in the parameterising torus $\mathbb{T}$ representing the pattern P_v, then

$$P_v + E = q(E + v) \subset q(NS) \subset \Omega\,. \tag{4.1}$$

Thus the non-singular points in the torus parameterisation give a large portion of Ω, and to obtain the full pattern space we need to complete this subset in the usual manner.

As noted, if v and v' are both non-singular points of $\mathcal{E}$, then the resulting patterns $\mathcal{P}_v$ and $\mathcal{P}_{v'}$ are locally equivalent, repetitive Delone sets, and hence share the same pattern space Ω. It suffices therefore to work with a single choice of non-singular v, and for simplicity of notation we shall assume from now on that $0 \in \mathcal{E}$ is non-singular and take $v = 0$; as in Definition 4.2 we shall denote the resulting pattern just by $\mathcal{P}$.

Recall that the space Ω may be considered as the completion of the set of all E-translates of $\mathcal{P}$ with respect to some appropriate *pattern metric*, such as the following. We may define a metric on subsets of E by declaring, for subsets $C, D \subset E$,

$$d(C,D) = \inf\left\{ \frac{1}{r+1} \;\middle|\; \begin{array}{l} \text{there exists } x, y \in B_{\frac{1}{r}} \text{ with} \\ \big(B_r \cap (C-x)\big) \cup \partial B_r = \big(B_r \cap (D-y)\big) \cup \partial B_r \end{array} \right\}$$

where B_r is the closed ball around 0 of radius r in E. In essence this metric is declaring C and D to be close if, up to a small translation, they are identical up to a long distance from the origin. The precise values or form of this metric will not be important for our discussion, indeed there are many possible variations, as discussed elsewhere, but rather the topology it generates.

As observed above, the set of all E-translates of $\mathcal{P}$ is given by

$$\mathcal{P} + E = \{\mathcal{P}_v \,|\, v \in E\}$$

and all these v are of course non-singular since 0 is. In terms of the torus parameterisation, Remark 4.7, this set corresponds to an irrationally sloped copy of E embedded in $\mathbb{T}$: the embedding property follows from the irrational position assumption made for the cut and project scheme, Definition 4.1(c). We arrive at our first two, equivalent descriptions of the pattern space Ω for $\mathcal{P}$, working along the lines indicated in Rermark 4.9.

Theorem 4.10.

1. *The pattern space Ω for $\mathcal{P}$ is given by the completion of $q(E) \subset \mathbb{T}$ with respect to the pattern metric.*
2. *The pattern space Ω for $\mathcal{P}$ is given by the completion of $q(NS) \subset \mathbb{T}$ with respect to the pattern metric.*

Proof. The first statement is really the definition of Ω under the assumption that q is one-to-one on E. The second statement follows from the inclusions (4.1). □

Remark 4.11. The reader may find it helpful to consider the following perspective. The set of patterns corresponding to points of $q(E)$ in the parameter space (which are also in 1 to 1 correspondence to the points $E \subset \mathcal{E}$) give all the translates of $\mathcal{P}$ in Ω. If $\mathcal{P}$ is obtained by placing the acceptance strip Σ at some point in $\mathcal{E}$, these can be conceived as the set of all patterns obtained by sliding the acceptance strip

along itself inside $\mathcal{E}$. The larger set of points, $q(NS)$, correspond to all patterns in Ω which may be obtained by placing the acceptance strip *somewhere* in $\mathcal{E}$, with the only constraint that no lattice point is allowed to lie on the boundary of Σ. What of the points corresponding to placing Σ in a position which *does* have a lattice point on its boundary, i.e., in a singular position? Such a positioning of Σ can be considered as a limit of a sequence of non-singular positions of the strip (or in fact of a sequence of points in $\mathcal{P} + E$ since $q(E)$ is dense in $q(NS)$), and so at first thought could correspond to a pattern that is a limit (i.e., a point in the completion of $q(E)$) of points in $\mathcal{P} + E$. However, a little more thought shows that the limit pattern actually depends on *how* the non-singular point is approached through singular points: if Σ has a lattice point, γ say, on its boundary, and we approximate this via a sequence of strip positions all of which have γ inside the strip, we will get a different limit pattern than if we approximate though a sequence which all have γ outside of the strip. The conclusion of this is that the points on $q(S)$, the image of the singular points in $\mathbb{T}$, correspond to multiple points in Ω. Moreover, the map $\mu\colon \Omega \to \mathbb{T}$, while it is 1 to 1 on the patterns corresponding to non-singular points, is multiple to 1 over the singular points.

5. A bundle structure on the pattern space Ω

In this section we develop the models for Ω of Section 4 so as to build descriptions of Ω as bundles over tori, and in due course be able to apply the ideas of Part I to obtain workable approaches to analysing $H^*(\Omega)$.

We start with the decomposition of $\mathcal{E}$ as $E \oplus F$. Let us define F_{NS} as $F \cap NS$, the set of non-singular points on the subspace F. As the E-translate of a non-singular point is non-singular, we have a decomposition of NS as the product $E \times F_{NS} \subset \mathcal{E}$. The lattice Γ of course acts on $\mathcal{E}$, but as it takes non-singular points to non-singular points, it also acts on $E \times F_{NS}$. In turn, it induces an action on F_{NS} by projection, as in Remark 4.4. Let us denote by F_c the completion of F_{NS} in the metric given by the sum of the pattern metric with the (usual) euclidean metric on F.[4] Following from the irrational position assumption of the projection scheme, the space F_c is totally disconnected; the closure (in F_c) of the interior of the acceptance domain K is a clopen set, and a basis of clopen sets is given by unions, intersections and complements of Γ-translates of this set. As in the case of F and of F_{NS}, the lattice Γ acts on F_c.

[4]This construction, which is crucial to the development of the models used to compute, and is integral to the discussion for the rest of this Part, can be thought of as follows. The subspace F_{NS} is obtained from F by removing the singular points; a Cauchy sequence in F_{NS} in the euclidean metric will converge to some, possibly singular, point of F and completion with the euclidean metric alone will just recover F again. However, two Cauchy sequences which converge in the euclidean metric to the same singular point could well converge to distinct points when we add in the pattern metric, depending on how they approach the singular point – exactly as in the discussion of Remark 4.11 above. As in that remark, there is a map $F_c \to F$ which is 1 to 1 on points in F_{NS}, and multiple to 1 over the singular points.

A diagrammatic view of all this is given by the following pullback diagram.

$$\begin{array}{ccc} \mathbb{Z}^N & & \mathbb{Z}^N \\ \downarrow & & \downarrow \\ E \times F_c & \longrightarrow & E \times F \\ \downarrow & & q \downarrow \\ \Omega & \xrightarrow{\mu} & \mathbb{T} \end{array}$$

Here $\mathbb{Z}^N \to E \times F \to \mathbb{T}$ is the fibration representing the quotient $\mathbb{T} = \mathcal{E}/\mathbb{Z}^N$ and the left-hand column is the pullback via the map μ. In short, we have Ω as the quotient $(E \times F_c)/\mathbb{Z}^N$.

Remark 5.1. The direct realisation of Ω as a bundle over a d-torus now proceeds as follows. As Γ acts on F_c by translation, we can split Γ as $\mathbb{Z}^d \oplus \mathbb{Z}^n$ where the $\mathbb{Z}^n$ acts freely on F_c. If we denote $F_c/\mathbb{Z}^n$ by X_c, a compact set (in fact a Cantor set) which still carries the remaining $\mathbb{Z}^d$ action, we obtain, using the relevant projection actions, the description

$$\Omega = (E \times X_c)/\mathbb{Z}^d \,.$$

The map $X_c \to (\texttt{point})$ now induces a map

$$\Omega = (E \times X_c)/\mathbb{Z}^d \to E/\mathbb{Z}^d = \mathbb{T}^d \,,$$

giving Ω the structure of a bundle over the d-torus with fibre X_c. As in Part I we obtain a relation

$$H^*(\Omega) = H^*(\mathbb{Z}^d; C(X_c; \mathbb{Z})) = H_{d-*}(\mathbb{Z}^d; C(X_c; \mathbb{Z})) \tag{5.1}$$

where the second equality comes from Poincaré duality (Appendix, Theorem 9.10).

It will turn out more convenient to pursue the *homology* version in this equation, in order not to get the indices too complicated: this is a purely book-keeping simplification, and maybe not one that seems that plausible to the reader given we have already had to supply the homology index as $d - *$, but bear with us, its convenience will be explained more later on.

However, in practice, neither of these equalities turns out to be exactly the most convenient way to analyse $H^*(\Omega)$ from the geometric data of the projection scheme. Instead, first note that the group homology of the equality (5.1) may be rewritten, using Remark 9.8

$$H^*(\Omega) = H_{d-*}(\mathbb{Z}^d; C(X_c; \mathbb{Z})) = H_{d-*}\left(\mathbb{Z}^d \oplus \mathbb{Z}^n; C(X_c; \mathbb{Z}) \otimes \mathbb{Z}[\mathbb{Z}^n]\right) \,.$$

We may identify $C(X_c; \mathbb{Z}) \otimes \mathbb{Z}[\mathbb{Z}^n]$ with $C_o(F_c; \mathbb{Z})$, the compactly supported, continuous, integer-valued functions on F_c. While philosophically this description is to be preferred as it removes the apparently arbitrary choice of splitting of $\mathbb{Z}^N$, it has the further practical advantage that the space F_c is somewhat easier to

describe than X_c, a description we shall turn to in a moment. The following summarises our main computational tool, one that is essentially reflecting the description $\Omega = (E \times F_c)/\mathbb{Z}^N$.

Theorem 5.2.

$$H^*(\Omega) = H_{d-*}(\mathbb{Z}^N; C_o(F_c; \mathbb{Z}))\,.$$

Let us turn now to the question of identifying F_c and the function space $C_o(F_c; \mathbb{Z})$. We approach $C_o(F_c; \mathbb{Z})$ through F_{NS} and the function space $C_o(F_{NS}; \mathbb{Z})$. Here, we consider the space F_{NS} as the complement in F, with the euclidean topology, of the singular points in F; these latter are simply the points

$$\Gamma + \partial K,$$

the Γ orbit (in F) of the boundary of the acceptance domain. As the Γ orbit of any point in F is dense in F, it is immediate from this that the space F_{NS} is also totally disconnected. We would like to identify the $\mathbb{Z}\Gamma$-algebras $C_o(F_c; \mathbb{Z})$ and $C_o(F_{NS}; \mathbb{Z})$; there is however a topological subtlety to be addressed first.

The space $C_o(F_{NS}; \mathbb{Z}))$ can be thought of as the compactly supported, integer-valued functions on F with discontinuities only at points of the subspace $\Gamma + \partial K$. In [7, 11] we introduced the notion of a $\mathcal{C}$*-tope*, a compact, polyhedral region of F with boundary contained in this singular set $\Gamma + \partial K$. The Γ action on F, which takes singular points to singular points, clearly takes $\mathcal{C}$-topes to $\mathcal{C}$-topes, and this realises the Γ action on $C_o(F_{NS}; \mathbb{Z}))$. On the other hand, $C_o(F_c; \mathbb{Z}))$ can be identified with the algebra generated by indicator functions on certain specific examples of $\mathcal{C}$-topes, namely those given by unions and intersections of Γ-translates of the actual acceptance domain K and its complement: why should this be the same algebra? In practice it need not be, though for large classes of examples these two algebras do agree, and even when they don't there is still much that can be said (see Remark 5.9 below). For now let us make the definition

Definition 5.3. Say that a projection scheme is *topially correct* if every $\mathcal{C}$-tope can be constructed via a finite sequence of unions, intersections and complements of polyhedra of the form $K + \pi_F(\gamma)$, for $\gamma \in \Gamma$.

If a scheme is topially correct then the algebras $C_o(F_c; \mathbb{Z})$ and $C_o(F_{NS}; \mathbb{Z})$ will agree and we can describe the algebra $C_o(F_c; \mathbb{Z})$ purely in terms of $\mathcal{C}$-topes. We return below to when this property holds, but its advantage is that the algebra of all $\mathcal{C}$-topes admits in many standard cases a good description, and one we will exploit in the next section.

In principle, the information needed to understand the algebra of $\mathcal{C}$-topes is completely described given the data defining the particular projection scheme, but let us introduce a couple of further assumptions on the projection schemes we will consider. These will allow for a very tractable understanding of this set and so of the coefficient module needed by Theorem 5.2 to describe $H^*(\Omega)$. Schemes satisfying the first assumption, Definition 5.4, receive a traditional name; it is nominally the case studied in [7], though the analysis there actually covers a much wider class of patterns, as we note below.

Definition 5.4. Suppose we take for $\mathcal{E}$ the vector space $\mathbb{R}^N$, and the lattice Γ as the corresponding integer lattice $\mathbb{Z}^N \subset \mathbb{R}^N$. Suppose the acceptance domain K is the π_F image of any translate of a unit cube I^N of $\mathbb{R}^N$, and that the translate is such that K is in non-singular position. Then the resulting pattern is known as a *canonical* cut and project pattern.

Lemma 5.5 ([7]). *Any canonical cut and project pattern is topially correct.*

Lemma 5.6. *Suppose we have a canonical cut and project pattern. Then for each face f_i of K, the set $f_i + \pi_F(\Gamma)$ contains the affine space spanned by f_i.*

Proof. This follows, for example, by noting that as K is the π_F-projection of some unit cube $Q = I^N + v$, the Γ-translates $Q + \Gamma$ tessellate $\mathbb{R}^N$, and in particular, the boundary faces of Q lie in the affine $(N-1)$-dimensional spaces they span. □

This lemma presents a property crucial to the accessible computation of cohomology developed in [7] and is worthy of abstraction, as found in [11]. In fact, the methods of [7] are immediately applicable to any cut and project pattern satisfying the following definition, Lemma 5.6 being one of the basic properties upon which the constructions of [7] are based.

Definition 5.7. We call a cut and project scheme *almost canonical* if it is topially correct and if for each face f_α of K, the set $f_\alpha + \pi_F(\Gamma)$ contains the affine space spanned by f_α.

Of course any canonical pattern is almost canonical in this sense by virtue of Lemmas 5.5 and 5.6, but the class of almost canonical patterns was first formally introduced by Julien in [14] in his study of the asymptotic complexity of cut and project patterns and is certainly much larger.

Remark 5.8. Apart from allowing us to identify the algebras $C_o(F_c;\mathbb{Z})$ and $C_o(F_{NS};\mathbb{Z})$, the crucial point of the *almost canonical* assumption is that the subspace $\Gamma + \partial K$ (and hence the algebra of $\mathcal{C}$-topes) can be described more simply via a collection of affine hyperplanes in F. Specifically, for each face f_α of the acceptance domain K, let W_α be the $(n-1)$-dimensional affine subspace spanned by f_α. Then F_{NS} is the complement of the Γ orbits of the set of W_α. Of course two or more W_α's may be in the same Γ orbit as each other. We let I_{n-1} be an indexing set for a complete (but non-redundant) set of W_α's, with just one representative from each Γ orbit. Then

$$F_{NS} = F \setminus \left(\Gamma + \bigcup_{\alpha \in I_{n-1}} W_\alpha \right) \quad \text{and} \quad NS = \mathcal{E} \setminus \left(E + \Gamma + \bigcup_{\alpha \in I_{n-1}} W_\alpha \right).$$

We will use this in the next section to develop a systematic approach to computing $H^*(\mathbb{Z}^N; C_o(F_c;\mathbb{Z}))$ via the geometry and combinatorics of the W_α and their intersections.

Remark 5.9. It remains to be asked, when will a cut and project scheme that satisfies the analogue of Lemma 5.6, that is, each set $f_\alpha + \pi_F(\Gamma)$ containing the affine space spanned by f_α, also be topially correct? It can be shown that one sufficient condition for this is that the subgroups of Γ which stabilise all the various *singular spaces* (see next section) act densely on their respective spaces, a condition readily checked in all the standard examples. Moreover, when this condition does *not* hold, it may also be shown that the resulting cohomology $H^*(\Omega)$ is then infinitely generated, and so in such cases rather obviates the need for the calculational machinery anyway.

6. Filtrations and exact sequences

Our aim is to compute, for an almost canonical pattern, the group homology $H_*(\mathbb{Z}^N; C_o(F_c;\mathbb{Z}))$, being identified with the Čech cohomology $H^{d-*}(\Omega)$. Our main tool is to place the coefficients in this group homology, $C_o(F_c;\mathbb{Z})$, into an exact sequence whose other terms are more simply related to the geometry of the affine hyperplanes W_α. From a homological point of view, we shall replace $C_o(F_c;\mathbb{Z})$ by a resolution, obtained from the natural dimensional filtration of F given by the singular set $\Gamma + \bigcup_{\alpha\in I_{n-1}} W_\alpha$.

Intersections of some of the Γ-translates of the W_α may be empty, or they may be affine subspaces of F of smaller dimension. We shall call all the translates of the W_α's, and all the lower-dimensional intersections, by the common name *singular spaces*. Let $\mathcal{P}_r$ denote the set of all singular spaces of dimension r (so, $\mathcal{P}_{n-1}$ is just the set of all $\gamma + W_\alpha$).

Definition 6.1. Let C_r, for $0 \leqslant r < n$ denote the $\mathbb{Z}\Gamma$ module generated by indicator functions on r-dimensional facets of $\mathcal{C}$-topes. An element of C_r may be thought of as a (finite sum of) compactly supported, integer-valued functions on a singular r-space, with discontinuities at lower-dimensional singular subspaces. We denote $C_o(F_c;\mathbb{Z})$ by C_n.

Theorem 6.2. *For an almost canonical pattern, there is an exact sequence of* $\mathbb{Z}\Gamma$*-modules and* Γ*-equivariant boundary maps*

$$0 \to C_n \xrightarrow{\delta} C_{n-1} \xrightarrow{\delta} \cdots \xrightarrow{\delta} C_0 \xrightarrow{\epsilon} \mathbb{Z} \to 0, \tag{6.1}$$

where δ *is given by the cellular boundary map on* $\mathcal{C}$*-topes and* ϵ *the augmentation map defined as follows. The module* C_0 *is generated by indicator functions on zero-dimensional singular spaces; denote such a function by* 1_p *for some* $p \in \mathcal{P}_0$*. Then* ϵ *is given by* $\epsilon(1_p) = 1$*.*

Proof. Consider the singular set $\Gamma + \bigcup_{\alpha\in I_{n-1}} W_\alpha$ as a limit of collections of these hyperplanes, where at each stage we take only a locally finite number of the $\gamma + W_\alpha$. Considered as a 'decoration' of $F = \mathbb{R}^n$, this can be thought of as a cellular decomposition of F, and as such there is an associated cellular chain complex,

similar to the sequence (6.1): its exactness and the final copy of $\mathbb{Z}$ correspond to the fact that the cellular homology of $\mathbb{R}^n$ is $\mathbb{Z}$ in degree 0 and 0 otherwise. The exact sequence (6.1) now follows by passing to the limit. □

This is a powerful computational tool, though it is perhaps best illustrated by working some examples. Before that, let us sketch two of the ways in which such a state of affairs may be utilised.

The first way, and the one followed in [7, 6] and the initial way used in [11] is to split (6.1) into a set of short exact sequences of $\mathbb{Z}\Gamma$ modules

$$\begin{aligned} &0 \to C_0^0 \to C_0 \to \mathbb{Z} \to 0\,, \\ &0 \to C_1^0 \to C_1 \to C_0^0 \to 0\,, \\ &\qquad \cdots \\ &0 \to C_n \to C_{n-1} \to C_{n-2}^0 \to 0\,. \end{aligned} \tag{6.2}$$

where C_r^0 denotes the kernel of the homomorphism $C_r \xrightarrow{\delta} C_{r-1}$. As noted in the Appendix (Theorem 9.6), a short exact sequence $\mathbb{Z}\Gamma$ modules gives a long exact sequence in group homology. Then the set (6.2) gives an inductive process of working up from knowledge of the singular 0-spaces, the singular 1-spaces, and so on, to a final calculation of $H_*(\Gamma; C_n)$. Clearly the complexity of this operation increases with the codimension, n, though there is relatively little increase in complexity for fixed n with an increase in the dimension parameter d, a fact in contrast to, for example, the standard approaches to computing cohomology for substitution patterns.

The final exact sequence of this family runs

$$\cdots \to H_{*+1}(\Gamma; C_{n-2}^0) \to H_*(\Gamma; C_n) \to H_*(\Gamma; C_{n-1}) \to H_*(\Gamma; C_{n-2}^0) \to \cdots \tag{6.3}$$

and gives the last part of a computation for $H^{d-*}(\Omega) = H_*(\Gamma; C_n)$; indeed, for codimension 2 patterns, it is the principal stage of the computation.

The full detailed application of this approach is provided in [11] Sections 5 and 6 (see also [7] Chapter V), and we shall illustrate a particular example of this method in the next section, but let us also note a second approach to computation that can be gained from the sequence (6.1).

The sequence (6.1) can be viewed as having three main parts: mentally bracket the sequence

$$0 \to C_n \xrightarrow{\delta} \left[C_{n-1} \xrightarrow{\delta} \cdots \xrightarrow{\delta} C_0 \right] \xrightarrow{\epsilon} \mathbb{Z} \to 0\,. \tag{6.4}$$

Moving to a slightly wider perspective of homological algebra (detailed, for example, in [11]) this gives rise to a single exact sequence

$$\cdots \to H_{*+n}(\Gamma; \underline{A}_*) \xrightarrow{j_*} H_{*+n}(\Gamma; \underline{T}_*) \xrightarrow{m_*} H_*(\Gamma; C_n) \to H_{*+n-1}(\Gamma; \underline{A}_*) \to \cdots\,; \tag{6.5}$$

see [11] Section 3.3 for definitions of the terms involved, but the relevant fact is that $H_{*+n}(\Gamma; \underline{T}_*)$ may be identified with the homology of the N-torus, with a shift

in dimension, and the other term, $H_{*+n}(\Gamma; \underline{A}_*)$, a group associated to the middle part of the sequence (6.4), specifically, the total homology of the bigraded complex

$$\underline{P}_* \otimes_{\mathbb{Z}\Gamma} \left[C_{n-1} \xrightarrow{\delta} \cdots \xrightarrow{\delta} C_0 \right]$$

for $\underline{P}_*$ some projective $\mathbb{Z}\Gamma$ resolution of $\mathbb{Z}$. Under a further assumption on the projection schemes considered (the *rationality assumptions* of [11] Section 4) this sequence can be identified as the long exact sequence

$$\begin{array}{ccccccccc} \cdots \to & H_r(\mathbb{A}) & \xrightarrow{\alpha_*} & H_r(\mathbb{T}) & \xrightarrow{m_*} & H_{r-n}(\Gamma; C_n) & \longrightarrow & H_{r-1}(\mathbb{A}) & \to \cdots \\ & & & \| & & \| & & & \\ & & & H^{N-r}(\mathbb{T}) & \xrightarrow{\mu^*} & H^{N-r}(\Omega) & & & \end{array}$$

where $\mu\colon \Omega \to \mathbb{T}$ is the quotient map of Remark 4.11, and $\mathbb{A}$ is a certain subspace of $\mathbb{T}$ defined in terms of the projection data (specifically, a union of subtori, one for each element of I_{n-1}). Moreover, this is the same sequence as Kalugin constructs, via a rather different approach, in [15].

Not surprisingly, given their common origin, the long exact sequences (6.3) and (6.5) are related, and a 'ladder' of exact sequence can be made from them allowing both to be utilised together in computations. An example of this can be found in the discussion of codimension 3 computations in [11].

Part III: Doing the computations

7. Ammann–Beenker

In this section we employ the ideas of Part II to work through the computations of a particular example. The one we choose is that of Ammann–Beenker [2], a two-dimensional tiling with codimension 2. The results for this example have already appeared several times in the literature, see for example [11, 12], but the computations here are presented in an annotated form for the novice calculator, and we hope will contribute to the accessibility of these types of calculation as well as shedding light on the theory of Part II above.

As can be seen from [11] many more examples have already been computed, including ones of both higher dimension and codimension. As noted earlier, the higher the codimension, the more complex the computation, and the interested reader will find a full discussion in [11] of the sort of treatment needed for a codimension 3 example, however many of the essential ideas are covered in the codimension 2 case. In the next section we make some notes on the sort of phenomena and problems that have to be faced in more general situation.

The Ammann–Beenker (sometimes also known as the *Octagonal* tiling) is, like the Penrose tilings, a pattern that can be described as a canonical projection pattern, a substitution or via matching rules. Here we sketch its description using a projection scheme.

Consider the following scheme. We take $\mathcal{E} = \mathbb{R}^4$ with $\Gamma = \mathbb{Z}^4 \subset \mathbb{R}^4$ the integer lattice. Let v_α for $\alpha = 1, \ldots, 4$ be the four unit vectors

$$(1,0,0,0) \qquad (0,1,0,0) \qquad (0,0,1,0) \qquad (0,0,0,1)$$

which both generate Γ and form a basis for $\mathcal{E}$. Consider the linear map $\mathbb{R}^4 \to \mathbb{R}^4$ given with respect to this basis by the matrix

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -1 & 0 & 0 & 0 \end{pmatrix}.$$

This is a rotation of order 8 and has two two-dimensional eigenplanes, one where the action is rotation by $\pi/4$, the other by $3\pi/4$; take the former for E and the latter for F. The acceptance domain K is given by the projection to F of the standard unit cube I^4 in $\mathbb{R}^4$. The special symmetries of this example mean that K is a perfect octagon.

As for any canonical pattern, the singular points $\Gamma + \partial K$ form complete affine hyperplanes in F – of course in this case, of codimension 2, that means one-dimensional lines. As in Remark 5.8 this singular set can be described as the Γ orbit of a finite set of lines, here it will be 4 lines corresponding to the four directions of the sides of the acceptance domain, which in turn are the projections of the four unit vectors v_α. Explicitly, we let W_α, $\alpha = 1, \ldots, 4$, be the one-dimensional subspace of F spanned by $\pi_F(v_\alpha)$. The W_α form four rotationally symmetric lines in F with $W_{\alpha+1}$ the rotation of W_α through $\frac{\pi}{4}$. The full singular set in F may be envisaged (perhaps after enough beer) as the subset of F given by taking the projected lattice $\pi_F(\Gamma)$, a dense subset of F, and for each point of this lattice placing a copy of the four singular lines passing though that point.

It will be the geometry and combinatorics of this arrangement of lines and their intersections that gives the necessary information to allow us to compute, so a few further details should be given. We will not prove our assertions, but they can all be readily checked with the aid of a few simple diagrams, which we encourage the reader to sketch.

First, as noted before, Γ acts on F by translation, and, as the singular lines are precisely the Γ orbits of the W_α, this action takes singular lines to singular lines. As $\Gamma = \mathbb{Z}^4$, we will use the v_α, $\alpha = 1, \ldots, 4$ also to denote a set of generators of Γ. It will be useful to have a description of the stabiliser subgroups of the singular lines; these are as follows. The stabiliser of each W_α is of rank 2, and specifically the stabilisers are given by

$$\Gamma^{W_1} = \langle v_1, v_2 - v_4 \rangle, \quad \Gamma^{W_2} = \langle v_2, v_1 + v_3 \rangle,$$
$$\Gamma^{W_3} = \langle v_3, v_2 + v_4 \rangle, \quad \Gamma^{W_4} = \langle v_4, v_1 - v_3 \rangle.$$

Next, we consider the intersections of the lines. By an *intersection point* we mean a point in F where two or more of the lines $\gamma + W_\alpha$ cross. Clearly (by construction) there are the lattice points $\pi_F(\Gamma)$ at each of which there are

precisely four singular lines meeting. There are other intersection points as well, where precisely two singular lines meet. Examples of these are

$$(v_1 + W_4) \cap W_2 \qquad \text{and} \qquad (v_2 + W_3) \cap W_1 \,.$$

These three examples of intersection points lie in distinct Γ orbits, and the three resulting orbits contain all intersection points. The Γ action on the intersection points is clearly free.

We are now ready to begin computing. We use the method of short exact sequences sketched after Theorem 6.2. Here there are just two sequences

$$\begin{aligned} 0 \to C_0^0 \to C_0 \to \mathbb{Z} \to 0\,, \\ 0 \to C_2 \to C_1 \to C_0^0 \to 0\,. \end{aligned} \tag{7.1}$$

The first gives rise to a long exact sequence in group homology

$$\cdots \to H_r(\Gamma; C_0) \to H_r(\Gamma; \mathbb{Z}) \to H_{r-1}(\Gamma; C_0^0) \to H_{r-1}(\Gamma; C_0) \to \cdots \,. \tag{7.2}$$

We take the terms in turn. The coefficient module C_0 is the $\mathbb{Z}\Gamma$ module of indicator functions on the intersection points: we have already noted that there are three distinct orbits of intersection points, and the Γ action on them is free. So C_0 is the free $\mathbb{Z}\Gamma$ module of rank 3. In particular, $H_0(\Gamma; C_0)$ is then $\mathbb{Z}^3$, one copy of $\mathbb{Z}$ for each orbit, and, by Remark 9.7, $H_r(\Gamma; C_0)$ vanishes for $r > 0$.

The next term, $H_r(\Gamma; \mathbb{Z})$ is simply the group homology of Γ with constant coefficients $\mathbb{Z}$. As discussed towards the end of the Appendix, this is equivalent to the homology of the 4-torus, and indeed it will be useful to write this homology in terms of the exterior power operation:

$$H_r(\Gamma; \mathbb{Z}) = H_r(\mathbb{T}^4) = \Lambda^r(\mathbb{Z}^4) = \mathbb{Z}^{\binom{4}{r}} \,.$$

Putting this together gives us a computation for the third term, $H_*(\Gamma; C_0^0)$, which we can then use as part of the computation of $H^r(\Omega) = H_{2-r}(\Gamma; C_2)$ via the second sequence in (7.1). The vanishing of $H_r(\Gamma; C_0)$ for $r > 0$ gives isomorphisms

$$H_r(\Gamma; C_0^0) = H_{r+1}(\Gamma; \mathbb{Z}) = \mathbb{Z}^{\binom{4}{r+1}} \qquad \text{for } r \geqslant 1\,,$$

and in low dimensions an exact sequence

$$0 \to H_1(\Gamma; \mathbb{Z}) \to H_0(\Gamma; C_0^0) \to H_0(\Gamma; C_0) \to H_0(\Gamma; \mathbb{Z}) \to 0,$$

i.e.,

$$0 \to \mathbb{Z}^4 \to H_0(\Gamma; C_0^0) \to \mathbb{Z}^3 \to \mathbb{Z} \to 0\,.$$

Basic algebra (in particular using the fact that $\mathbb{Z}$ is projective) solves this sequence and we obtain

$$H_r(\Gamma; C_0^0) = \begin{cases} 0 & \text{if } r > 3 \\ \mathbb{Z} & \text{if } r = 3 \\ \mathbb{Z}^4 & \text{if } r = 2 \\ \mathbb{Z}^6 & \text{if } r = 1 \\ \mathbb{Z}^6 & \text{if } r = 0. \end{cases} \tag{7.3}$$

We now turn to the second long exact sequence, that arising from the second line of (7.1). This runs

$$\cdots \to H_r(\Gamma; C_1) \to H_r(\Gamma; C_0^0) \to H_{r-1}(\Gamma; C_2) \to H_{r-1}(\Gamma; C_1) \to \cdots . \qquad (7.4)$$

We have a computation for $H_r(\Gamma; C_0^0)$, but we need also a computation for $H_r(\Gamma; C_1)$. To do this we consider geometrically what the module C_1 is representing. This module is the set of (compactly supported, integer-valued) functions on the singular lines $\gamma + W_\alpha$, with discontinuities at the intersection points on these lines. As such we can decompose C_1, first as

$$C_1 = \bigoplus_{i=1}^{4} \{\text{Functions on lines parallel to } W_\alpha\}$$

and this allows us to decompose the homology as

$$H_r(\Gamma; C_1) = \bigoplus_{i=1}^{4} H_r\big(\Gamma; \{\text{Functions on lines parallel to } W_\alpha\}\big) .$$

We can simplify this still further: the lines parallel to W_α are given by $\gamma + W_\alpha$ as γ ranges over Γ/Γ^{W_α}, thus Γ/Γ^{W_α} acts freely on the module of functions on lines parallel to W_α and we can write

$$\{\text{Functions on lines parallel to } W_\alpha\} = C_1(W_\alpha; \mathbb{Z})) \otimes_{\mathbb{Z}} \mathbb{Z}[\Gamma/\Gamma^{W_\alpha}]$$

where $C_1(W_\alpha; \mathbb{Z}))$ denotes the compactly supported, integer-valued functions on the *single line* W_α with discontinuities at the intersection points on that line. Hence, by Remark 9.8, we have

$$H_r\left(\Gamma; \{\text{Functions on lines parallel to } W_\alpha\}\right) = H_r(\Gamma^{W_\alpha}; C_1(W_\alpha; \mathbb{Z})) .$$

Recall from our discussion of intersection points earlier that each stabiliser Γ^{W_α} is of rank 2, and each line W_α has precisely two distinct orbit classes of intersection point.

We can compute $H_r(\Gamma^{W_\alpha}; C_1(W_\alpha; \mathbb{Z}))$ in much the same way as we computed $H_r(\Gamma^{W_\alpha}; C_0^0)$: as in the complex in Theorem 6.2, there is an exact sequence

$$0 \to C_1(W_\alpha; \mathbb{Z}) \xrightarrow{\delta} C_0(W_\alpha; \mathbb{Z}) \xrightarrow{\epsilon} \mathbb{Z} \to 0,$$

where $C_0(W_\alpha; \mathbb{Z})$ denotes the functions on the intersection points on W_α. We get an induced long exact sequence in group homology[5], now taking homology for the group Γ^{W_α}. We obtain

$$\begin{aligned} \cdots \quad & \to H_r(\Gamma^{W_\alpha}; C_0(W_\alpha; \mathbb{Z})) \to H_r(\Gamma^{W_\alpha}; \mathbb{Z}) \to H_{r-1}(\Gamma^{W_\alpha}; C_1(W_\alpha; \mathbb{Z})) \to \\ & \to H_{r-1}(\Gamma^{W_\alpha}; C_0(W_\alpha; \mathbb{Z})) \to \cdots . \end{aligned}$$

[5]It is here that it pays to work in homology, not cohomology: if we were to perform these computations for the lower-dimensional parts of the singular set in group cohomology then the indices would become a mess. The situation would become still worse in higher codimensions where there were multiple values of r for which we would have to carry out similar operations on modules C_r with $n > r > 0$.

Just as in the earlier case, $H_r(\Gamma^{W_\alpha}; C_0(W_\alpha;\mathbb{Z}))$ vanishes for $r > 0$, is $\mathbb{Z}^2$ for $r = 0$ (there are two orbits of intersection points on each W_α) and $H_r(\Gamma^{W_\alpha};\mathbb{Z})$ is essentially the homology of the 2-torus (as Γ^{W_α} is of rank 2). We conclude

$$H_r(\Gamma^{W_\alpha}; C_1(W_\alpha;\mathbb{Z})) = \begin{cases} 0 & \text{if } r > 1 \\ \mathbb{Z} & \text{if } r = 1 \\ \mathbb{Z}^3 & \text{if } r = 0 \end{cases} \tag{7.5}$$

and hence

$$H_r(\Gamma; C_1) = \bigoplus_{i=1}^{4} H_r(\Gamma^{W_\alpha}; C_1(W_\alpha;\mathbb{Z})) = \begin{cases} 0 & \text{if } r > 1 \\ \mathbb{Z}^4 & \text{if } r = 1 \\ \mathbb{Z}^{12} & \text{if } r = 0. \end{cases} \tag{7.6}$$

We turn to the final computation, for $H_*(\Gamma; C_2)$ from the short exact sequence $0 \to C_2 \to C_1 \to C_0^0 \to 0$. The fact that $H_r(\Gamma; C_1) = 0$ for $r > 1$ yields an isomorphism $H_r(\Gamma; C_2) = H_{r+1}(\Gamma; C_0^0)$ in this range, but this just gives $H_r(\Gamma; C_2) = \mathbb{Z}$ for $r = 2$ and vanishes for higher r. The more problematic part is the lower end of the sequence, which runs

$$\begin{aligned} 0 = H_2(\Gamma; C_1) &\to H_2(\Gamma; C_0^0) \to \\ &\to H_1(\Gamma; C_2) \to H_1(\Gamma; C_1) \to H_1(\Gamma; C_0^0) \to \\ &\to H_0(\Gamma; C_2) \to H_0(\Gamma; C_1) \to H_0(\Gamma; C_0^0) \to 0 \end{aligned}$$

which by the computations above becomes

$$0 \to \mathbb{Z}^4 \to H_1(\Gamma; C_2) \to \mathbb{Z}^4 \xrightarrow{\rho} \mathbb{Z}^6 \to H_0(\Gamma; C_2) \to \mathbb{Z}^{12} \to \mathbb{Z}^6 \to 0\,. \tag{7.7}$$

This sequence does not have a unique solution without further knowledge[6] of the homomorphism ρ.

The homomorphism $\rho\colon H_1(\Gamma; C_1) \to H_1(\Gamma; C_0^0)$ can be identified by retracing the analysis of these particular homology groups above. It can be shown that there is a commutative diagram

$$\begin{array}{ccccc} \Lambda^2\mathbb{Z}^4 & = & H_2(\Gamma;\mathbb{Z}) & \xrightarrow{\cong} & H_1(\Gamma; C_0^0) \\ \uparrow{\scriptstyle i_*} & & \uparrow{\scriptstyle i_*} & & \uparrow{\scriptstyle \rho} \\ \bigoplus_{i=1}^4 \Lambda^2\mathbb{Z}^2 & = & \bigoplus_{i=1}^4 H_2(\Gamma^{W_\alpha};\mathbb{Z}) & \xrightarrow{\cong} & \bigoplus_{i=1}^4 H_1(\Gamma^{W_\alpha}; C_1(W_\alpha;\mathbb{Z})) \\ & & & & \| \\ & & & & H_1(\Gamma; C_1) \end{array}$$

where the right-hand horizontal maps marked $\cong$ are the homomorphisms in the sequence (7.2) and the corresponding one for $H_*(\Gamma^{W_\alpha}; C_1(W_\alpha;\mathbb{Z}))$. The homomorphism marked i_* is that induced from the inclusions $\Gamma^{W_\alpha} \subset \Gamma$. Thus i_* and hence ρ are completely computable, and the kernel is a rank 1 subgroup of $\mathbb{Z}^4$.

[6] We'll pick up on this further in the next section, but one might note that this sequence **does** determine the *Euler characteristic* – the alternating sum of the ranks of the groups for $H^*(\Omega)$: changing the rank of ρ changes the ranks of H_1 and H_0 equally, keeping the sum $\operatorname{rk} H_0(\Gamma; C_2) - \operatorname{rk} H_1(\Gamma; C_2) + \operatorname{rk} H_2(\Gamma; C_2)$ equal at 5.

We can now finalise the calculation from the exact sequence (7.4): the groups involved are all free abelian and there are no extension problems to solve. We obtain

Theorem 7.1. *The Čech cohomology of the pattern space for the Ammann–Beenker projection scheme is*

$$H^r(\Omega) = H_{2-r}(\Gamma; C_2) = \begin{cases} 0 & \text{if } r > 2 \\ \mathbb{Z}^9 & \text{if } r = 2 \\ \mathbb{Z}^5 & \text{if } r = 1 \\ \mathbb{Z} & \text{if } r = 0. \end{cases}$$

8. Further remarks

We conclude with some short remarks on further topics concerning the cohomology of projection patterns.

8.1. Finite generation

We omit a detailed discussion of what sort of answers are likely to arise as the cohomology of a projection pattern, but there are some simple things that can be said about whether or not the cohomology groups that arise are finite or infinitely generated. After all, by identifying $H^*(\Omega)$ in terms of the group homology of the infinitely generated module C_n, it is not *a priori* clear that the answers will be finitely generated.

The first result to note is

Theorem 8.1 ([7, IV.6.7]). *A necessary condition for $H^*(\Omega)$ to be finitely generated is that n divides N.*

Indeed, the number $\nu = \frac{N}{n} = 1 + \frac{d}{n}$ is an important constant related to the order of stabiliser groups of the singular spaces in the full analysis of [11]. For now we note that the requirement that n divides N, and hence d, means that the only projection patterns with finitely generated cohomology of dimension 2 are those of codimension 1 or 2, and for $d = 3$ the only patterns will be of codimension 1 or 3. The cohomology of patterns as defined in Definition 4.1 of codimension 1 are always finitely generated.

While this is a necessary condition, apart from the case of codimension 1, it is not sufficient, indeed in some sense almost every (almost canonical) projection pattern of codimension $n > 1$ will have infinitely generated cohomology. A complete condition is given by the following. We denote by L_0 the number of Γ orbits of intersection points in the singular set of F.

Theorem 8.2 ([7, IV.2.9, V.2.4]). *$H^*(\Omega) \otimes \mathbb{Q}$ is a finite dimension, graded vector space if and only if L_0 is finite.*

It is this result that lies behind the final observation in Remark 5.9. Also, it may be easily checked, for example, see [13], that a projection scheme that does not satisfy the conclusion of Lemma 5.6 will fail to have L_0 finite.

8.2. Torsion

For rather trivial reasons, the cohomology of any tiling of dimension 1 is free of torsion. To see this consider, for example, the fact that any tiling space of dimension d may be written as an inverse limit

$$\Omega = \lim_{\leftarrow} \{\cdots \to X_r \to X_{r-1} \to \cdots \to X_1\}$$

for d-dimensional CW complexes X_r. Then $H^*(\Omega)$ is the direct limit of the $H^*(X_r)$. In the case $d = 1$ the spaces X_r are, up to homotopy equivalence, one point unions of circles, and so each $H^*(X_r)$ is free abelian; the direct limit of free abelian groups cannot contain torsion.

For higher dimensions, there may be torsion. The 'smallest' examples occur with $n = d = 2$ and include, for example, the Tübingen Triangle Tiling [1, 16]; see [9, 11] for a discussion. For these parameters torsion may appear in $H^2(\Omega) = H_0(\mathbb{Z}^4; C_2)$, and it arises through the analogue of the map ρ in (7.7) having torsion in its cokernel. Computationally this is straightforward – it just comes out in the computations and is no more problematic than computing a dimension 2 codimension 2 example without torsion. The story is rather different in higher codimensional examples (see next subsection).

As of this writing, it is not clear that the existence of torsion in the cohomology of a tiling space has any very significant implications or geometric 'meaning'.

8.3. The complications of higher codimensions

The codimension 2 example detailed in Section 7 demonstrated all the essential ingredients needed to be considered for the computation of any codimension 2 projection scheme. For $n = 2$ but higher dimension d there is more to be done, more cohomology groups to compute, but the method goes over and no significantly new ideas are needed (this is not to underestimate however the computational complexity of sorting out, for example, the singular spaces, their stabiliser groups and their intersections).

Moving up just one more codimension, to $n = 3$, introduces significant complications; these fall into two types. The fundamental complex (6.1) now gives rise to three short exact sequences of coefficients; the first two can be addressed using techniques similar to those used for $n = 2$, but by the time we consider the final resulting long exact sequence, there are problems. Both types of problems relate to the solving of exact sequences. In the codimension 2 case we already saw that it was necessary to use further geometric information to compute the action of the homomorphism we labelled ρ; in the more complex cases there are similar homomorphisms that need describing, but which no longer have such simple formulations in terms of the underlying geometric data. The second type of problem concerns the solution of extension problems. A short exact sequence of abelian groups

$$0 \to A \to B \to C \to 0 \tag{8.1}$$

has unique solution $B = A \oplus C$ if C is free abelian: this was a fact we used several times in the codimension 2 calculation. By the time we compute codimension 3 examples, exact sequences such as (8.1) occur in which C can contain torsion; here the solution is no longer unique, and further arguments are needed to establish which of the possible solutions is the correct one. The reader will glimpse some of the issues from the discussion of the integral cohomology calculations of [11] §6.

Of course if we limit ourselves to computing *rational* cohomology, these extension problems go away, and even the first problems become tractable. In [11] Theorem 6.5 we give a complete formulaic answer to $H^*(\Omega; \mathbb{Q})$ for any dimension 3, codimension 3 almost canonical projection pattern.

8.4. Euler characteristic

Finally, we saw in the footnote to Section 7 that the Euler characteristic of the Ammann–Beenker tiling could be easily read off our machinery, without even pursuing it to the final cohomology calculations. This is true in general (and for higher codimension 3 as well, assuming one can successfully describe all the singular subspaces and their stabiliser groups). The diligent reader who follows the calculations of Section 7 will have no problem proving the following result. The number L_0 is the number of Γ orbits of intersection points in the singular set in F, while L_0^α denotes the number of Γ^{W_α} orbits of intersection points in W_α.

Theorem 8.3. *The Euler characteristic χ of a dimension 2, codimension 2 almost canonical projection pattern is $(\sum_i L_0^\alpha) - L_0$. In particular $\chi \geqslant 1$.*

The observation about χ being positive follows from the fact that an intersection point has to be the intersection of at least two of the W_α's, so each orbit contributes 1 to L_0, but it also contributes 1 to at least *two* of the L_0^α's.

The smallest possible value of χ, namely 1, may indeed be realised. Consider the following variation on the Ammann–Beenker example [13]. We keep the same E, F and lattice Γ in $\mathcal{E} = \mathbb{R}^4$, but replace the acceptance domain K by the parallelogram with sides $\pi_F(v_1)$ and $\pi_F(v_2)$. This is not a canonical pattern, but it is still almost canonical. Then there is precisely one orbit of intersection points, so $L_0 = L_0^1 = L_0^2 = 1$, giving $\chi = 2 - 1 = 1$.

Appendix

9. Notes on Čech and group cohomology

9.1. Čech cohomology

This is not a formal introduction and definition of Čech cohomology, rather it is what is needed for the job in hand. We assume throughout that the space X is compact, and of course are essentially thinking of either Cantor sets or tiling space like objects. The most essential facts are

Theorem 9.1.

1. *On any finite CW complex, Čech cohomology, singular cohomology and cellular cohomology all agree.*
2. *The Čech cohomology* (*unlike the other two above*) *satisfies the* continuity axiom*, that if the space X is given as an inverse limit*
$$X = \lim_{\leftarrow} \left\{ \cdots \to X_i \to \cdots X_2 \to X_1 \to X_0 \right\}$$
then the Čech cohomology of X is the direct limit of the cohomologies of the X_n,
$$H^n(X) = \lim_{\rightarrow} \left\{ H^n(X_0) \to H^n(X_1) \to \cdots \to H^n(X_i) \to \cdots \right\}.$$

Corollary 9.2. *If X is a totally disconnected space, then the Čech cohomology $H^n(X)$ vanishes for $n > 0$ and $H^0(X)$ may be identified with the continuous $\mathbb{Z}$-valued functions $C(X;\mathbb{Z})$ on X.*

Proof. Any such space may be written as an inverse limit $X = \lim_{\leftarrow} X_i$ of finite CW complexes X_i with the homotopy type of a finite set of points. The continuity axiom and standard calculations for discrete spaces immediately tell us that $H^n(X)$ must vanish for $n > 0$; the result for degree 0 follows from the identification of $H^0(Y)$as $C(Y;\mathbb{Z})$ for any finite discrete space Y, and the definition of the inverse limit topology on X. □

Remark 9.3. A note is in order for when considering, for example, rational Čech cohomology. A moment's thought shows that $C(X;\mathbb{Q})$ and $C(X;\mathbb{Z}) \otimes \mathbb{Q}$ are not the same: if X contains an infinite number of points then the latter will contain functions taking an infinite number of distinct values, while the latter will not. Following the argument of the corollary above, we see that in this case, the correct identification is
$$H^0(X;\mathbb{Q}) = C(X;\mathbb{Z}) \otimes \mathbb{Q}.$$

9.2. Group cohomology

We collect here the details and properties of group homology and cohomology which we have used in the main sections of this article. Necessarily, this can only be a brief tour through the facts, and the reader interested in seeing more detail or proofs of the results below should consult one of the principal texts. Good sources that cover this material, from various viewpoints, include [3, 4, 20, 21].

The perspective that is particularly close to the motivation for group cohomological techniques we have sketched in this article is that of groups of invariants and coinvariants, and their corresponding derived functors. We begin with introducing these notions. In what follows, a *G-module M* refers to an abelian group which is a module over the group-ring $\mathbb{Z}G$.[7]

[7]Recall, the group-ring $\mathbb{Z}G$ is the free $\mathbb{Z}$-module on basis elements $\{g \,|\, g \in G\}$, and the ring structure is generated by declaring the product of two basis elements, g_1 and g_2 say, to be the basis element corresponding to the group product $g_1 g_2$; this extends linearly. Thus a typical

Definition 9.4. Suppose G is a group and M a G-module. Define the *subgroup of G-invariants* as

$$M^G = \{m \in M \mid gm = m\} = \hom_{\mathbb{Z}G}(\mathbb{Z}; M)$$

where $\mathbb{Z}$ carries the trivial G-action, and the group of *coinvariants* as the quotient

$$M_G = M/\langle m - gm\rangle = \mathbb{Z} \otimes_{\mathbb{Z}G} M,$$

i.e., the quotient of M by the submodule generated by all elements of the form $m - gm$ as m and g run over all elements of M and G respectively.

Considered as functors from the category of G-modules to abelian groups, the invariants and coinvariants are *half-exact*, meaning that if $0 \to M \to N \to Q \to 0$ is a short exact sequence of G modules and G-maps, the sequences

$$0 \to M^G \to N^G \to Q^G \qquad \text{and} \qquad M_G \to N_G \to Q_G \to 0$$

are both exact, but $N^G \to Q^G$ is not necessarily surjective, and $M_G \to N_G$ is not necessarily injective. Group cohomology and homology can be seen as a remedy for this: indeed, with the addition of certain normalising assumptions, this can be taken as the defining property of each, see, for example, [20] for an elegant introduction along these lines.

Theorem 9.5. *For each group G and $n = 0, 1, 2, \ldots$ there are functors $H^n(G; -)$ from G-modules to abelian groups satisfying the properties*

(i) *$H^0(G; M) = M^G$, the G-invariants of M, and*
(ii) *if $0 \to M \to N \to Q \to 0$ is a short exact sequence of G modules and G-maps, then there is a long exact sequence of abelian groups*

$$\begin{aligned} 0 \to H^0(G; M) \to H^0(G; N) \to H^0(G; Q) \to H^1(G; M) \to \cdots \\ \cdots \to H^n(G; M) \to H^n(G; N) \to H^n(G; Q) \to H^{n+1}(G; M) \to \cdots . \end{aligned}$$

Similarly, for homology, we have

Theorem 9.6. *For each group G and $n = 0, 1, 2, \ldots$ there are functors $H_n(G; -)$ from G-modules to abelian groups satisfying the properties*

(i) *$H_0(G; M) = M_G$, the G-coinvariants of M, and*
(ii) *if $0 \to M \to N \to Q \to 0$ is a short exact sequence of G modules and G-maps, then there is a long exact sequence of abelian groups*

$$\begin{aligned} \cdots \to H_n(G; M) \to H_n(G; N) \to H_n(G; Q) \to H_{n-1}(G; M) \to \cdots \\ \cdots \to H_1(G; M) \to H_0(G; M) \to H_0(G; N) \to H_0(G; Q) \to 0 . \end{aligned}$$

element of $\mathbb{Z}G$ is a sum of the form $\sum_{g_i \in G} n_i g_i$, where the $n_i \in \mathbb{Z}$ and only a finite number of them are non-zero.

There are various approaches to establishing the existence of the functors $H^*(G;-)$ and $H_*(G;-)$, perhaps the most familiar being via projective resolutions. We suppose that $\underline{P}_*$ is a projective resolution of $\mathbb{Z}$ as $\mathbb{Z}G$ modules, that is, an exact sequence of $\mathbb{Z}G$ modules and $\mathbb{Z}G$ module maps

$$\cdots \to P_n \to P_{n-1} \to \cdots \to P_2 \to P_1 \to P_0 \to \mathbb{Z} \to 0\,.$$

(Recall that a module is projective if and only if it is a direct summand of a free module; the free modules are thus themselves projective. In practice, in the situations we are interested in, it is usually more convenient to consider *free* resolutions, where the P_n are free $\mathbb{Z}G$ modules.)

We may then construct the homology and cohomology groups as the homology of the associated complexes

$$\begin{aligned} &H^*(G;M) = \text{Homology of} \\ &\quad \cdots \leftarrow \hom_{\mathbb{Z}G}(P_n, M) \leftarrow \hom_{\mathbb{Z}G}(P_{n-1}, M) \leftarrow \cdots \leftarrow \hom_{\mathbb{Z}G}(P_0, M); \end{aligned}$$

$$\begin{aligned} &H_*(G;M) = \text{Homology of} \\ &\quad \cdots \to P_n \otimes_{\mathbb{Z}G} M \to P_{n-1} \otimes_{\mathbb{Z}G} M \to \cdots \to P_0 \otimes_{\mathbb{Z}G} M. \end{aligned}$$

It is precisely the fact noted above that the invariant and coinvariant functors (i.e., the functors $\hom_{\mathbb{Z}G}(-,M)$ and $-\otimes_{\mathbb{Z}G} M$) do not take exact sequences to exact sequences that makes these homologies non-trivial.

Remark 9.7. It may be easily checked from this construction that if M is itself a *free* (or even projective) $\mathbb{Z}G$ module, then these functors *are* exact and in particular $H_n(G;M)$ vanishes for all $n > 0$.

Remark 9.8. An extension of this observation provides an important simplification in many of our calculations for the cut and project tilings. Suppose G decomposes as $G = K \times G/K$ for some summand K, and the G module M takes the form $N \otimes \mathbb{Z}[G/K]$ where G/K acts trivially on the K-module N and G acts in the obvious way on $\mathbb{Z}[G/K]$, (i.e., M is the free $\mathbb{Z}[G/K]$ module induced from the K-module N). Then there is an equivalence

$$H_*(G;M) \cong H_*(K;N)\,.$$

Example 9.9. There are many ways to construct a projective resolution of $\mathbb{Z}$ by $\mathbb{Z}G$ modules, and for computational purposes the smaller and simpler the better. We consider a straightforward, but relevant, family of examples, that of the free abelian groups $G = \mathbb{Z}^d$.

(a) First let us consider the case where G is the group of integers $\mathbb{Z}$. A particularly convenient $\mathbb{Z}\mathbb{Z}$ resolution, here in fact by *free* $\mathbb{Z}\mathbb{Z}$ modules, of $\mathbb{Z}$ is given by the sequence

$$0 \to \mathbb{Z}\mathbb{Z} \xrightarrow{1-s} \mathbb{Z}\mathbb{Z} \xrightarrow{\epsilon} \mathbb{Z} \to 0\,. \tag{9.1}$$

We explain the homomorphisms involved. Recall that $\mathbb{Z}\mathbb{Z}$ is a free $\mathbb{Z}$ module on a countable generating set indexed by the elements of $\mathbb{Z}$; let us denote the mth

such basis element by **m**. Then s is the shift map, sending **m** to **m+1**, and the 'augmentation' homomorphism $\epsilon\colon \mathbb{ZZ} \to \mathbb{Z}$ acts sending each **m** to 1.

As the resolution we have been able to use contains non-trivial modules P_n only for $n = 0$ and 1, and $P_n = 0$ for all $n \geqslant 2$, it is immediate that $H^n(\mathbb{Z}; M)$ and $H_n(\mathbb{Z}; M)$ vanish for all $n \geqslant 2$, for any M.

Computing, say the homology with coefficients M the constant module $\mathbb{Z}$, in the remaining dimensions proceeds as follows. Tensoring the resolution above by $- \otimes_{\mathbb{ZZ}} \mathbb{Z}$ yields the chain complex

$$0 \to \mathbb{ZZ} \otimes_{\mathbb{ZZ}} \mathbb{Z} = \mathbb{Z} \xrightarrow{(1-s)\otimes 1} \mathbb{ZZ} \otimes_{\mathbb{ZZ}} \mathbb{Z} = \mathbb{Z} \to 0$$

and it is easy to check that the 'boundary' homomorphism $(1 - s) \otimes 1$ is zero, yielding $H_0(\mathbb{Z}; \mathbb{Z}) = H_1(\mathbb{Z}; \mathbb{Z}) = \mathbb{Z}$. The calculation for cohomology is similar.

(b) To handle the case of $G = \mathbb{Z}^d$ for $d > 1$, we note that resolutions of products of groups can be given by (tensor) products of resolutions of the individual groups. In particular, we may take products of d copies of the resolution (9.1) above to obtain a free $\mathbb{ZZ}^d$ resolution of $\mathbb{Z}$, obtaining a sequence where the n^{th} module P_n is of the form

$$\mathbb{ZZ}^d \oplus \overset{\binom{d}{n}}{\cdots} \oplus \mathbb{ZZ}^d\,. \tag{9.2}$$

In principal this can be used to compute: as in the case $d = 1$, all the boundary homomorphisms after tensoring with $- \otimes_{\mathbb{ZZ}^d} \mathbb{Z}$ vanish and the homology and cohomology may be simply idenifitied: in the case of homology,

$$H_n(\mathbb{Z}^n; \mathbb{Z}) = \left(\mathbb{ZZ}^d \oplus \overset{\binom{d}{n}}{\cdots} \oplus \mathbb{ZZ}^d \right) \otimes_{\mathbb{ZZ}^d} \mathbb{Z} = \mathbb{Z} \oplus \overset{\binom{d}{n}}{\cdots} \oplus \mathbb{Z}\,.$$

(In general, homology and cohomology of products of groups may be deduced from the application of Künneth theorems which in this case are particularly simple taking the (co)homology of the product of copies of $\mathbb{Z}$ to the tensor product of copies of the (co)homology of a single $\mathbb{Z}$; other groups can give more complicated relationships.) However, the next observation links this calculation to a better known computation.

There is an important topological analogue of the algebraic construction of a resolution. For each group G there is a topological space, the *classifying space of* G, denoted BG, enjoying various properties. In fact, just as there are many possible projective resolutions for any given G, so too are there many possible spaces which are perfectly good candidates for being called BG: nevertheless, they are all homotopy equivalent to each other, and the space BG is only defined up to homotopy equivalence.

What the space BG classifies does not concern us here, but two facts are pertinent. The first is that for any group G and G module M there are (natural) equivalences

$$H^n(G; M) \cong H^n(BG; M) \qquad \text{and} \qquad H_n(G; M) \cong H_n(BG; M)$$

the right-hand object in each equivalence being (singular) (co)homology with coefficients in the module M.

The second point to bear in mind is that there can be drawn very close parallels between cell complex structures of possible BG's and free $\mathbb{Z}G$ resolutions (in fact, this is one way in which the equivalences just noted may be demonstrated). Given a free $\mathbb{Z}G$ resolution of $\mathbb{Z}$, that is one where the n^{th} group P_n is a direct sum of copies of $\mathbb{Z}G$, a cell complex may be constructed (the *geometric realisation of* $\underline{P}_*$) which has one n-cell for each $\mathbb{Z}G$ summand in P_n; similarly the boundary maps from the n-cells to the $n-1$ cells are determined by the data coded in the homomorphism $P_n \to P_{n-1}$ (again, see texts such as [3, 20] for details).

In the case of $G = \mathbb{Z}$ and the resolution (9.1) above, we construct a space $B\mathbb{Z}$ using just one 0-cell, and one 1-cell, whose ends must of course be both attached to the single 0-cell: we obtain a copy of the circle $\mathbb{S}^1$. Similarly, taking products as sketched above, we find that a model of $B(\mathbb{Z}^d)$ is given by the d-torus $\mathbb{T}^d = \mathbb{S}^1 \times \overset{d}{\cdots} \times \mathbb{S}^1$.

The final important fact about group cohomology which we use in this article is that the finitely generated, free abelian groups $\mathbb{Z}^d$ are so-called *Poincaré duality* groups.[8] This means that they satisfy

Theorem 9.10. *For any $\mathbb{Z}^d$ module M, there are isomorphisms*

$$H^r(\mathbb{Z}^d; M) \cong H_{d-r}(\mathbb{Z}^d; M)\,.$$

Note that in particular, since H^0 and H_0 are respectively the invariant and coinvariant functions, we may identify the *top*-dimensional *cohomology* $H^d(\mathbb{Z}^d; M)$ with the coinvariants M_G, and the top dimensional homology $H_d(\mathbb{Z}^d; M)$ with the invariants M^G. We used this fact in Section 2.

Acknowledgement

The author is indebted to Alan Forrest who asked the questions that gave rise to the ideas of using group cohomology in this subject, and to Alan Forrest, Franz Gähler and Johannes Kellendonk with whom most of the main work here has been coauthored. Thanks should also be recorded to Ian Putnam and Claire Irving with whom the author has had may useful discussions on this topic.

References

[1] M. Baake, P. Kramer, M. Schlottmann, and D. Zeidler, Planar patterns with fivefold symmetry as sections of periodic structure in 4-space, *Int. J. Mod. Phys. B* 4 (1990) 2217–2268.

[2] F.P.M. Beenker, Eindhoven University of Technology Report 82-WSK-04 (1982).

[3] D.J. Benson, Representations and Cohomology: Volume 2, Cohomology of Groups and Modules, Cambridge studies in advanced mathematics, CUP, 1998.

[8] This is not unrelated to Poincaré duality for manifolds, and can for instance be deduced from the fact that the classifying spaces $B\mathbb{Z}^d$ may be realised as manifolds, viz the tori $\mathbb{T}^d$ as above.

[4] K. Brown, Cohomology of Groups, Springer-Verlag, 1982.
[5] A.H. Forrest and J.R. Hunton, The cohomology and K-theory of commuting homeomorphisms of the Cantor set, *Ergod. Th. and Dynam. Sys.*, 19 (1999) 611–625.
[6] A.H. Forrest, J.R. Hunton and J. Kellendonk, Cohomology of canonical projection tilings, *Comm. Math. Phys.*, 226 (2002) 289–322.
[7] A.H. Forrest, J.R. Hunton and J. Kellendonk, Topological invariants for projection method patterns, *Mem. Amer. Math. Soc.* 758 (2002).
[8] F. Gähler, Lectures given at workshops *Applications of Topology to Physics and Biology*, Max-Planck-Institut für Physik komplexer Systeme, Dresden, June 2002, and *Aperiodic Order, Dynamical Systems, Operator Algebras and Topology*, Victoria, British Columbia, August, 2002.
[9] F. Gähler, Lectures and unpublished notes on torsion in the homology of the Tübingen Triangle Tiling, Banff and elsewhere, (2004).
[10] F. Gähler, J.R. Hunton and J. Kellendonk, Integer Čech cohomology of icosahedral projection tilings, *Z. Kristallogr.* 223 (2008) 801–804.
[11] F. Gähler, J.R. Hunton and J. Kellendonk, Integral cohomology of rational projection method patterns, *Algebraic and Geometric Topology*, 13 (2013) 1661–1708.
[12] F. Gähler and J. Kellendonk, Cohomology groups for projection tilings of codimension 2, *Materials Science and Engineering.* 438–440 (2000) 294–296.
[13] C. Irving, Euler characteristics and cohomology for quasiperiodic projection patterns, PhD thesis, University of Leicester, 2006.
[14] A. Julien, Complexity and cohomology for cut-and-projection tilings, *Ergodic Theory and Dynamical Systems*, 30 (2010) 489–523.
[15] P. Kalugin, Cohomology of quasiperiodic patterns and matching rules, *J. Phys. A: Math. Gen.* 38 (2005) 3115–3132.
[16] R. Klitzing, M. Schlottmann, and M. Baake, Perfect matching rules for undecorated triangular tilings with 10-, 12-, and 8-fold symmetry, *Int. J. Mod. Phys. B* 7 (1993) 1455–1473.
[17] R. Moody, Model sets: A Survey, From Quasicrystals to More Complex Systems, eds. F. Axel, F. Dénoyer, J.P. Gazeau, Centre de physique Les Houches, Springer-Verlag, 2000.
[18] I.F. Putnam, Non-commutative methods for the K-theory of C_*-algebras of aperiodic patterns from cut-and-project systems, *Comm. Math. Phys.*, 294(2010), 703–729.
[19] L. Sadun and R.F. Williams, Tiling spaces are Cantor set bundles, *Ergodic Theory and Dynamical Systems*, 23 (2003), 307–316.
[20] C.B. Thomas, Characteristic classes and the cohomology of finite groups, Cambridge Studies in Advanced Mathematics, CUP, 1986.
[21] C. Wiebel, An introduction to homological algebra, *Cambridge Studies in Advanced Mathematics* 38 (1994), CUP.

John Hunton
Department of Mathematical Sciences, Durham University
Science Laboratories, South Road
Durham, DH1 3LE, UK
e-mail: john.hunton@durham.ac.uk
URL: http://www.dur.ac.uk/john.hunton

Progress in Mathematics, Vol. 309, 137–194
© 2015 Springer Basel

Equicontinuous Factors, Proximality and Ellis Semigroup for Delone Sets

Jean-Baptiste Aujogue, Marcy Barge,
Johannes Kellendonk and Daniel Lenz

Abstract. We discuss the application of various concepts from the theory of topological dynamical systems to Delone sets and tilings. In particular, we consider the maximal equicontinuous factor of a Delone dynamical system, the proximality relation and the enveloping semigroup of such systems.

Mathematics Subject Classification (2010). 37B50, 52C22.

Keywords. Tiling, Delone set, maximal equicontinuous factor, dynamical spectrum, proximal relation, Ellis semigroup.

1. Introduction

Recurrence properties of patterns in a Euclidean point set may be effectively studied by means of an associated dynamical system. The idea, due originally to Dan Rudolph ([66]) and analogous to the notion of a sub-shift in Symbolic Dynamics, is as follows: Given a point set $\mathcal{L} \subset \mathbb{R}^N$, let $\Omega_\mathcal{L}$, the *hull* of $\mathcal{L}$, be the collection of all point sets in $\mathbb{R}^N$, each of which is locally indistinguishable from $\mathcal{L}$. There is a natural *local topology* on $\Omega_\mathcal{L}$, $\mathbb{R}^N$ acts on $\Omega_\mathcal{L}$ by translation, and the structure of $\mathcal{L}$ is encoded in the topological dynamics of the system $(\Omega_\mathcal{L}, \mathbb{R}^N)$. In this article we consider the correspondence between properties of the point set $\mathcal{L}$ and the closeness to equicontinuity of the system $(\Omega_\mathcal{L}, \mathbb{R}^N)$. (For an excellent survey of the dynamical properties of the hull of a point set, with an emphasis on sets arising from substitutions, see [65].)

Consider, for example, the 'crystalline' case of a completely periodic set $\mathcal{L} = \{a + kv_i : i = 1, \dots, N,\, k \in \mathbb{Z}\} \subset \mathbb{R}^N$ with $a \in \mathbb{R}^N$ and linearly independent $v_i \in \mathbb{R}^N$. The hull of $\mathcal{L}$ is in this case just the set of translates $\Omega_\mathcal{L} = \{\mathcal{L} - t : t \in \mathbb{R}^N\}$. Letting L be the lattice spanned by the basis $\{v_i\}_{i=1}^N$, we may identify $\Omega_\mathcal{L}$ with the torus $\mathbb{R}^N/L$ by $\mathcal{L} - t \leftrightarrow L - t$, and we see that the dynamical system $(\Omega_\mathcal{L}, \mathbb{R}^N)$ is simply the translation action on a compact abelian group. This is an *equicontinuous*

action, that is, the collection of homeomorphisms $\{\alpha_t : \mathcal{L}' \mapsto \mathcal{L}' - t\}_{t\in\mathbb{R}^n}$ of $\Omega_\mathcal{L}$ is an equicontinuous family; in particular, if $\mathcal{L}', \mathcal{L}'' \in \Omega_\mathcal{L}$ are *proximal* (that is, the distance between $\mathcal{L}' - t$ and $\mathcal{L}'' - t$ is not bounded away from zero) then $\mathcal{L}' = \mathcal{L}''$.

We are of course more interested in $\mathcal{L}$ that are highly structured (in this article, always Delone with finite local complexity, usually repetitive, and often Meyer) but not periodic. The hull will no longer be a group, and there will be distinct elements which are proximal so the action is no longer equicontinuous, but we can begin to understand the system $(\Omega_\mathcal{L}, \mathbb{R}^N)$, and thus the structure of $\mathcal{L}$, by comparing $(\Omega_\mathcal{L}, \mathbb{R}^N)$ with its largest (maximal) equicontinuous factor $(\Omega_{\max}, \mathbb{R}^N)$. Existence of a maximal equicontinuous factor is immediate: Let $R_{\max}$ be the intersection of all closed invariant equivalence relations R on $\Omega_\mathcal{L}$ for which the action on $\Omega_\mathcal{L}/R$ is equicontinuous. Then $\Omega_{\max} := \Omega_\mathcal{L}/R_{\max}$ is a maximal equicontinuous factor; uniqueness of $\Omega_{\max}$ (up to topological conjugacy) is a consequence of maximality. But understanding what is collapsed in passing from $\Omega_\mathcal{L}$ to $\Omega_{\max}$ requires a more concrete formulation of the equicontinuous structure relation $R_{\max}$. It is clear that elements $\mathcal{L}', \mathcal{L}'' \in \Omega_\mathcal{L}$ which are proximal must be identified in $\Omega_{\max}$, and for many of the familiar $\mathcal{L}$, $R_{\max}$ is just the proximal relation. Clearly, the equicontinuous structure relation is a closed equivalence relation; in general, however, the proximal relation is neither closed nor transitive. To illustrate this, consider the example[1] with $\mathcal{L} := -2\mathbb{N}\cup\mathbb{N} \subset \mathbb{R}$. The hull $\Omega_\mathcal{L}$ is the disjoint union of three $\mathbb{R}$-orbits, namely it contains besides the translates of $\mathcal{L}$ also those of $\mathcal{L}^+ := \mathbb{Z}$ and $\mathcal{L}^- := 2\mathbb{Z}$. It thus consists of two circles $\{\mathcal{L}^+ - t : t \in \mathbb{R}\} = \mathbb{R}/\mathbb{Z}$, $\{\mathcal{L}^- - t : t \in \mathbb{R}\} = \mathbb{R}/2\mathbb{Z}$ and the curve $\{\mathcal{L} - t : t \in \mathbb{R}\}$ winding from one to the other. Notice that the distance between $\mathcal{L} - t$ and $\mathcal{L}^+ - t$ goes to zero as $t \to +\infty$ and the distance between $\mathcal{L} - t$ and $\mathcal{L}^- - t$ goes to zero as $t \to -\infty$. Thus $\mathcal{L}$ is proximal with each of $\mathcal{L}^-$ and $\mathcal{L}^+$, but these latter are not proximal with each other. So the proximal relation is not transitive. Note also that, for $m \in \mathbb{Z}$, $\mathcal{L}^- = \mathcal{L}^- - 2m$ is proximal with $\mathcal{L} - 2m$ and the latter tends to $\mathcal{L}^+$ as $m \to \infty$. Thus the proximal relation is not topologically closed. The maximal equicontinuous factor is the circle $\Omega_{\max} = \mathbb{R}/\mathbb{Z}$ with ∞-to-1 factor map $\pi_{\max} : \Omega_\mathcal{L} \to \Omega_{\max}$ given by $\pi(\mathcal{L}' - t) = \mathbb{Z} - t$ for $\mathcal{L}' \in \{\mathcal{L}^-, \mathcal{L}, \mathcal{L}^+\}$.

In Section 3 we show that points of $\Omega_\mathcal{L}$ are identified in $\Omega_{\max}$ if and only if they have the same image under all continuous eigenfunctions of $(\Omega_\mathcal{L}, \mathbb{R}^N)$ (see Theorem 3.10) and in Section 4, several intrinsically defined variants of the proximal relation are discussed. These include a uniform version of proximality called *syndetic proximality* and *regional proximality*, which for minimal systems always agrees with the equicontinuous structure relation. Each of the proximal, syndetic proximal and regional proximal relations takes a 'strong' form for hulls of repetitive Meyer sets (repetitive Delone sets $\mathcal{L}$ for which $\mathcal{L} - \mathcal{L}$ is also Delone). For hulls of repetitive Delone sets, if the regional proximal relation is equal to a statistical version of proximality called *statistical coincidence*, all of these relations

[1]This example is not repetitive and hence is a bit misleading in its simplicity: the hull of any repetitive and non-periodic Delone set is considerably more complicated, having uncoutably many $\mathbb{R}^N$-orbits and being locally the product of a Cantor set with a Euclidean disk (see, for example, [67]).

are the same and the system is close to being equicontinuous in a specified sense (Corollary 4.15).

To measure the closeness of the system $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ to equicontinuity, one can consider the cardinality of a fiber $\pi_{\max}^{-1}(\xi)$, $\xi \in \Omega_{\max}$, of the equicontinuous factor map $\pi_{\max}$. But which ξ? By ergodicity of $(\Omega_{\max}, \mathbb{R}^N)$ there is $m \in \mathbb{N} \cup \{\infty\}$ such that $\sharp\pi_{\max}^{-1}(\xi) = m$ for Haar-almost all ξ, but there are other useful notions for the rank of $\pi_{\max}$. Two obvious ones are the *minimal rank* and *maximal rank*: $\mathrm{mr} := \inf_{\xi\in\Omega_{\max}} \sharp\pi_{\max}^{-1}(\xi)$ and $Mr := \sup_{\xi\in\Omega_{\max}} \sharp\pi_{\max}^{-1}(\xi)$. A third, which turns out to be extremely useful, is the *coincidence rank*, cr, defined as the supremum, over $\xi \in \Omega_{\max}$, of the supremum of cardinalities of subsets of $\pi_{\max}^{-1}(\xi)$ whose elements are pairwise non-proximal. For a repetitive Delone set $\mathcal{L}$ with hull $\Omega_{\mathcal{L}}$, the proximal relation and the equicontinuous structure relation $R_{\max}$ are the same if and only if $\mathrm{cr} = 1$ (Theorem 4.7). If $\mathrm{mr} < \infty$, then the proximality relation and $R_{\max}$ are the same if and only if the proximality relation is closed (Theorem 2.15 of [11]) and if $\mathrm{cr} < \infty$, then the proximal relation is closed if and only if $\mathrm{cr} = 1$ (Theorem 4.11). The meaning of the coincidence rank is revealed most clearly when $\Omega_{\mathcal{L}}$ is the hull of a repetitive Meyer set and the fiber distal points have full Haar measure: For R sufficiently large, for each $\xi \in \Omega_{\max}$, there is a set $A \subset \mathbb{R}^N$ of density 1 so that there are exactly cr distinct sets of the form $\mathcal{L}' \cap B_R(v)$, $\mathcal{L}' \in \pi_{\max}^{-1}(\xi)$, for each $v \in A$ (see Theorem 4.20). That is to say, viewed out to radius R, a fiber typically appears to have cardinality cr.

Due to its connection with pure point diffraction spectrum of $\mathcal{L}$ (see [17] and [46, 5]), the $\mathrm{mr} = 1$ case (that is, $\pi_{\max}$ is somewhere 1-1) has received the most attention. The following results (the first two due to Baake, Lenz, and Moody [6] and the third to Aujogue [2]) for the hull of a repetitive Meyer set $\mathcal{L}$ are discussed in Section 3.3:

- $\pi_{\max}$ is everywhere 1-1 if and only if $\mathcal{L}$ is completely periodic (Theorem 3.14);
- $\pi_{\max}$ is almost everywhere 1-1 if and only if $\mathcal{L}$ is a regular complete Meyer set (Theorem 3.15); and
- $\pi_{\max}$ is somewhere 1-1 if and only if $\mathcal{L}$ is a complete Meyer set (Theorem 3.16).

In the above, a *complete Meyer set* is a repetitive inter-model set whose window is the closure of its interior; such a set is *regular* if the boundary of the window has zero measure. This hierarchy gives a rather satisfying picture of the correspondence between injectivity properties of $\pi_{\max}$ and structural properties of the set $\mathcal{L}$.

When the coincidence rank is greater than 1, the equicontinuous structure relation is no longer given by proximality and the situation is considerably more complicated. We are able to make a few observations in Section 4 when the coincidence rank is known to be finite (as is the case for 'Pisot type' substitutive systems). For example, if $\mathrm{cr} < \infty$ for the dynamical system on the hull of an N-dimensional repetitive Delone set with finite local complexity, then the system is topologically conjugate to that on the hull of a repetitive Meyer set. Moreover, if the Delone set has no periods, then its topological eigenvalues are dense in $\hat{\mathbb{R}}^N$ (Theorem 4.13).

Also, for repetitive Delone sets whose system $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ has finite coincidence rank, if the set of points ξ in $\Omega_{\max}$ with the property that the points in the fiber $\pi_{\max}^{-1}(\xi)$ are pairwise non-proximal has full Haar measure, and if μ is any ergodic probability measure on $\Omega_{\mathcal{L}}$, then the continuous eigenfunctions generate $L^2(\Omega_{\mathcal{L}}, \mu)$ if and only if $\mathrm{cr} = 1$, and these conditions imply unique ergodicity of $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ (see Theorem 4.12). But a systematic understanding of the structure of $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ when the minimal rank is greater than 1 remains a challenging problem for the future.

One can view the hull of $\mathcal{L}$ and its maximal equicontinuous factor as compactifications of the acting group $\mathbb{R}^N$ (at least when $\mathcal{L}$ has no periods, so that $\mathbb{R}^N$ acts faithfully). Another compactification, which preserves more of the topology of $\Omega_{\mathcal{L}}$ than does $\Omega_{\max}$, while still introducing additional algebraic structure, is provided by the Ellis semigroup $E(\Omega_{\mathcal{L}}, \mathbb{R}^N)$. This is defined as the closure of the set of homeomorphisms $\{\alpha_t : \mathcal{L}' \mapsto \mathcal{L}' - t \,|\, t \in \mathbb{R}^n\} \subset \Omega_{\mathcal{L}}^{\Omega_{\mathcal{L}}}$ in the Tychonov topology with semigroup operation given by composition. For our considerations, two algebraic properties of the Ellis semigroup are particularly relevant: $E(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ has a unique minimal ideal if and only if the proximal relation is an equivalence relation; and two elements of $\Omega_{\mathcal{L}}$ are proximal if and only if they have the same image under some idempotent belonging to a minimal ideal of $E(\Omega_{\mathcal{L}}, \mathbb{R}^N)$. Consider the example above with $\mathcal{L} := -2\mathbb{N} \cup \mathbb{N}$. The sequence $(\alpha_{2m})_{m\in\mathbb{N}}$ converges to the element $\alpha^- \in E(\Omega_{\mathcal{L}}, \mathbb{R})$: α^- is the identity on each of the circles $\mathbb{T}^{\pm} := \{\mathcal{L}^{\pm} - t : t \in \mathbb{R}\}$ and collapses the curve $\{\mathcal{L} - t : t \in \mathbb{R}\}$ onto $\mathbb{T}^-$ by $\mathcal{L} - t \mapsto \mathcal{L}^- - t$. In particular, α^- is an idempotent ($\alpha^- \circ \alpha^- = \alpha^-$) and α^- identifies the proximal points $\mathcal{L}$ and $\mathcal{L}^-$. There is a similarly defined idempotent α^+ that identifies $\mathcal{L}$ and $\mathcal{L}^+$, and the Ellis semigroup is isomorphic with the disjoint union: $E(\Omega_{\mathcal{L}}, \mathbb{R}) \simeq \mathbb{T} \times \{\alpha^-\} \cup \mathbb{R} \times \{id\} \cup \mathbb{T} \times \{\alpha^+\}$, where $\mathbb{T} = \Omega_{\max}$ which is a group, and the operation is coordinate-wise (and non-abelian, since $\alpha^- \circ \alpha^+ = \alpha^-$ while $\alpha^+ \circ \alpha^- = \alpha^+$). There are two minimal (left) ideals, corresponding to $\mathbb{T} \times \{\alpha^+\}$ and $\mathbb{T} \times \{\alpha^-\}$, reflecting the fact that proximality is not transitive.

The Ellis semigroup is typically a very complicated gadget. One can't even expect countable neighborhood bases for the topology. It is thus surprising that there is a relatively simple algebraic and topological description of the Ellis semigroup of the hull of any *almost canonical* cut-and-project set, similar to that given in the previous paragraph. In this description, found by Aujogue in [2] and explained here in Section 5 (see also [3]), the finite submonoid of idempotents effectively captures the proximal structure of the hull. This family of examples suggests that the deep and well-developed abstract theory of the Ellis semigroup may find significant further application in the study of highly structured Delone sets.

We begin this article with a review of relevant facts from dynamics and basic constructions of Delone sets in Section 2. Maximal equicontinuous factors associated with Delone sets are considered in Section 3 and variants of the proximal relation are discussed in Section 4. The final Section 5 introduces the Ellis semigroup and concludes with the example of the semigroup of the hull of the Octagonal tiling.

2. Background

2.1. General background

In this article we consider actions of topological groups on topological spaces. The spaces, often denoted by X, will be compact Hausdorff spaces. The groups will be locally compact, σ-compact **abelian** groups and mostly denoted by G. The group composition will generally be written additively as $+$. The neutral element will be denoted by e. The dual of a group G consists of all continuous homomorphisms from G to the unit circle. It is equipped with the compact open topology, which means that a sequence of homomorphisms converges if it converges uniformly on every compact subset of G. The dual group is again a locally compact abelian group and will be denoted by $\widehat{G}$. As the elements of $\widehat{G}$ are maps on G there is a dual pairing between a group G and its dual group $\widehat{G}$. It will be denoted as $(\cdot, \cdot)$.

We will assume metrizability of X in some cases in order to ease the presentation of certain concepts. If the corresponding results are valid without the metrizability assumption we have stated them without this assumption. In certain cases we will also need the groups to be compactly generated. The main application we have in mind are Delone dynamical systems in Euclidean space. These systems are metrizable and the underlying group (Euclidean space) is compactly generated. Thus, **all** our results below apply in this situation.

2.2. Background on dynamical systems

Whenever the locally compact abelian group G acts on the compact space X by a continuous action

$$\alpha : \; G \times X \; \longrightarrow \; X \,, \quad (t, x) \; \mapsto \; \alpha_t x \,,$$

where $G \times X$ carries the product topology, the triple (X, G, α) is called a *topological dynamical system* over G. We will mostly suppress the action α in our notation and write

$$t \cdot x := \alpha_t x.$$

Accordingly, we will then also suppress α in the notation for the dynamical system and just write (X, G) instead of (X, G, α).

A dynamical system (X, G) is called *minimal* if, for all $x \in X$, the *G-orbit* $\{t \cdot x : t \in G\}$ is dense in X.

Let two topological dynamical systems (X, G) and (Y, G) over G be given. Then, a continuous map $\varrho : X \longrightarrow Y$ is called a *G-map* if $\varrho(t \cdot x) = t \cdot \varrho(x)$ holds for all $x \in X$ and $t \in G$. A G-map is called a *factor map* if it is onto. In this case, (Y, G) is called a *factor* of (X, G). Factor maps will mostly denoted by π. A G-map is called a *conjugacy* if it is a homeomorphism. Then, the dynamical systems are called *conjugate*. In this case, of course, each system is a factor of the other.

An important role in our subsequent considerations will be played by dynamical systems in which X is a compact group and the action is induced by a homomorphism into this group. In order to simplify the notation we provide a special name for such systems.

Definition 2.1 (Rotation on a compact abelian group). A dynamical system (X, G) is called a rotation on a compact abelian group if X is a compact abelian group and the action of G on X is induced by a homomorphism $j : G \longrightarrow X$ such that $t \cdot x = j(t) + x$ for all $t \in G$ and $x \in X$.

Remarks

- The name of rotation on a compact group comes from the example of a rotation by the angle $\alpha \in \mathbb{R}$ on the unit circle. In that case X is given by the unit circle $S^1 = \{z \in \mathbb{C} : |z| = 1\}$ and j is the map from the group $\mathbb{Z}$ of integers into S^1 mapping n to $e^{i\alpha n}$.
- If $G = \mathbb{R}$ such systems are also known as Kronecker flows. More generally, for G arbitrary they are known as Kronecker systems.
- We will think of the compact group as a form of torus. Accordingly, in the sequel we will often denote rotations on a compact group by $(\mathbb{T}, G)$.
- A rotation on a compact group is minimal if and only if j has dense range.

Whenever (X, G) is a dynamical system then $\chi \in \widehat{G}$ is called a *continuous eigenvalue* if there exists a continuous f on X with $f \neq 0$ and

$$f(t \cdot x) = (\chi, t) f(x)$$

for all $t \in G$ and $x \in X$. Such an f is then called an *continuous eigenfunction (to the eigenvalue χ).*

If (X, G) is minimal short arguments show the following: Firstly, any continuous eigenfunction has constant modulus and therefore does not vanish anywhere. Secondly, two continuous eigenfunctions to the same eigenvalue are linear dependent. In particular, the dimension of the space of all continuous eigenfunctions to a fixed eigenvalue is always one.

The set of all continuous eigenvalues is a subgroup of $\widehat{G}$. Indeed, the constant function 1 is always a continuous eigenfunction to the eigenvalue 0, the product of two continuous eigenfunction is a continuous eigenfunction (to the product of the eigenvalues) and the complex conjugate of a continuous eigenfunction is a continuous eigenfunction (to the inverse of the eigenvalue). The group of continuous eigenvalues of the dynamical system (X, G), equipped with the discrete topology, plays a crucial role in the subsequent considerations and we denote it by $\mathcal{E}_{\mathrm{top}}(X, G)$.

There is a strong connection between the group of continuous eigenvalues and a certain rotation on a compact group, called the *maximal equicontinuous factor*. This factor is at the heart of the investigations of this chapter. In fact, Section 3.3 is devoted to how the maximal equicontinuous factor controls the original dynamical system and Section 4 deals with the fine analysis of the equivalence relation induced by this factor.

For our investigations to be meaningful we will need non-triviality of the group of continuous eigenvalues. Dynamical systems without non-trivial continuous eigenvalues are called *topologically weakly mixing*. This property can be seen to be equivalent to transitivity of the product system with the diagonal action.

In the present description of dynamical systems we have so far been concerned with the topological point of view. Indeed, this is the main focus of our considerations. However, for certain issues we will need measure theoretical aspects as well.

As a compact space the set X underlying the dynamical system (X, G) carries naturally the Borel-σ-algebra which is the smallest σ-algebra containing all compact sets. A measure m on X is called *invariant* if $m(\alpha_t B) = m(B)$ for any Borel measurable set B and any $t \in G$. An invariant probability measure is a called *ergodic* if any Borel set B with $\alpha_t(B) = B$ for all $t \in G$ satisfies $m(B) = 0$ or $m(B) = 1$. Any dynamical system admits invariant probability measures and the ergodic measures can be shown to be the extremal points of the set of invariant probability measures (see, for example, the monographs [16, 76]). In particular, if a dynamical system admits only one invariant probability measure, then this measure is automatically ergodic. In this case the dynamical system is called *uniquely ergodic*.

Whenever m is an invariant probability measure on (X, G) the triple (X, G, m) is called a measure-preserving dynamical system. The action of G induces a unitary representation on $L^2(X, m)$ as follows: For any $t \in G$ there is a unitary map

$$T_t : L^2(X, m) \longrightarrow L^2(X, m),\ T_t f(x) = f(t \cdot x).$$

The behaviour of the action can be analysed through the behaviour of the representation T. This is sometimes known as *Koopmanism*.

As usual an element $f \in L^2(X, m)$ is called a *measurable eigenfunction* to the *measurable eigenvalue* $\chi \in \widehat{G}$ if

$$T_t f = (\chi, t) f$$

holds for all $t \in G$. Here, the equality is meant in the sense of L^2.

Ergodicity implies that the modulus of any measurable eigenfunction is constant almost surely and that two measurable eigenfunctions to the same eigenvalue are linearly dependent. The set of all measurable eigenvalues forms a subgroup of $\widehat{G}$. This subgroup, equipped with the discrete topology, is denoted by $\mathcal{E}_{meas}(X, T)$. If the closed subspace of $L^2(X, m)$ generated by the eigenfunctions agrees with $L^2(X, m)$ then (X, G, m) is said to have *pure point spectrum*.

2.3. Background on Delone sets

Let G be a locally compact abelian group. We will deal with subsets $\mathcal{L}$ of G. A subset $\mathcal{L}$ of G is called *uniformly discrete* if there exists an open neighborhood U of the identity in G such that

$$(x + U) \cap (y + U) = \emptyset$$

for all $x, y \in \mathcal{L}$ with $x \neq y$. A subset $\mathcal{L}$ of G is called *relatively dense* if there exists a compact neighborhood K of the identity of G such that

$$G = \bigcup_{x \in \mathcal{L}} (x + K).$$

A subset $\mathcal{L}$ of G is called a *Delone set* if it is both uniformly discrete and relatively dense. There is a natural action of G on the set $\mathcal{U}(G)$ of uniformly discrete sets in G via

$$G \times \mathcal{U}(G) \longrightarrow \mathcal{U}(G),\ (t, \mathcal{L}) \mapsto \mathcal{L} - t := \{x - t : x \in \mathcal{L}\}.$$

We refer to it as *translation action*.

Most prominent is the case $G = \mathbb{R}^N$. In that case one can express the above definitions using balls with respect to Euclidean distance. The open set U is expressed by an open ball and the compact set K by a closed ball. We denote the open ball with radius r around $x \in \mathbb{R}^N$ by $U_r(x)$ and the closed ball around x with radius R by $B_R(x)$. Then, $\mathcal{L} \subset \mathbb{R}^N$ is uniformly discrete if and only if there exists an $r > 0$ with $U_r(x) \cap U_r(y) = \emptyset$ for all $x, y \in \mathcal{L}$ with $x \neq y$, i.e., if and only if there exists an $r > 0$ such that the distance between any two different points of $\mathcal{L}$ is at least $2r$. Such a set will then be called *r-discrete*. The set $\mathcal{L} \subset \mathbb{R}^N$ is relatively dense if and only if there exists an $R > 0$ with $\mathbb{R}^N = \cup_{x \in \mathcal{L}} B_R(x)$, that is, if and only if any point of $\mathbb{R}^N$ has distance not exceeding R to $\mathcal{L}$.

Whenever $\mathcal{L}$ is a uniformly discrete subset of G a set of the form $(\mathcal{L} - x) \cap K$ with $x \in \mathcal{L}$ and K compact is called a *patch of* $\mathcal{L}$. A uniformly discrete subset $\mathcal{L}$ in G is said to have *finite local complexity* (FLC) if for any compact K in G the set

$$\{(\mathcal{L} - x) \cap K : x \in \mathcal{L}\}$$

is finite. This just means that there are only finitely many patches for fixed 'size' K. It is not hard to see that $\mathcal{L}$ has finite local complexity if and only if the set

$$\mathcal{L} - \mathcal{L} = \{x - y : x, y \in \mathcal{L}\}$$

is locally finite, i.e., has finite intersection with any compact subset of G. This in turn is equivalent to $\mathcal{L} - \mathcal{L}$ being closed and discrete. A Delone set with finite local complexity will be referred to as an *FLC Delone set*.

For an FLC Delone set $\mathcal{L}$ we define the *hull* $\Omega_{\mathcal{L}}$ to be the set of all Delone sets whose patches are also patches of $\mathcal{L}$. This set is obviously invariant under the translation action of G given above. Moreover, it is compact in a natural topology (discussed below). So, when equipped with the translation action, $\Omega_{\mathcal{L}}$ becomes a dynamical system, $(\Omega_{\mathcal{L}}, G)$, which we refer to as the *dynamical system associated to* $\mathcal{L}$.

When $G = \mathbb{R}^N$ it is possible to further characterize finite local complexity. A Delone set $\mathcal{L} \subset \mathbb{R}^N$ with $\mathbb{R}^N = \cup_{x \in \mathcal{L}} B_R(x)$ for some $R > 0$ has finite local complexity if and only if the set

$$\{(\mathcal{L} - x) \cap B_{2R}(0) : x \in \mathcal{L}\}$$

is finite [40]. Thus, in this case one needs to test for finiteness of the number of patches only for patches of a certain fixed size. For this reason, the hull of a Delone set with finite local complexity in $\mathbb{R}^N$ can be thought of as a geometric analogue to a subshift over a finite alphabet.

An *occurrence of the patch* $(\mathcal{L} - x) \cap K$ in a Delone set $\mathcal{L}$ is an element of

$$\{y \in \mathcal{L} : (\mathcal{L} - x) \cap K \subset (\mathcal{L} - y)\}.$$

A Delone set $\mathcal{L}$ is called *repetitive* if for any nonempty patch the set of occurrences is relatively dense. For an FLC Delone set $\mathcal{L}$, repetitivity is equivalent to minimality of the associated system $(\Omega_{\mathcal{L}}, G)$.

There are various equivalent approaches to define a topology on the set of all uniformly discrete sets. One is based on the identification of point sets with measures and the vague topology [13, 5], another uses uniform structures [5]. If $G = \mathbb{R}^N$ one can also make precise the idea of defining a metric by the principle that two sets are the ϵ-close if they coincide up to an error of ϵ on the $\frac{1}{\epsilon}$-ball around 0 [66, 71, 22]. In general the error of coincidence is measured with the help of the Hausdorff distance. A particularily elegant formulation of this idea uses the stereographic projection and has been worked out in detail in [52].

If the sets in question have finite local complexity, which is the only case which we consider in more detail here, then the description of the topology simplifies. A net [2] $(\mathcal{L}_\iota)_\iota$ in the hull $\Omega_{\mathcal{L}}$ of an FLC Delone set converges to $\mathcal{L}'$ if and only if there exists a net $(t_\iota)_\iota$ in G converging to e and for all compact $K \subset G$ an ι_K such that $(\mathcal{L}_\iota - t_\iota) \cap K = \mathcal{L}' \cap K$ for all $\iota > \iota_K$. Moreover, if $G = \mathbb{R}^N$ then the topology on hull $\Omega_{\mathcal{L}}$ of a Delone set with finite local complexity is induced by the metric [1]

$$d(\mathcal{L}, \mathcal{L}') = \inf \left\{ \frac{\epsilon}{\epsilon + 1} : \exists t, t' \in B_\epsilon(0) : B_{\frac{1}{\epsilon}}[\mathcal{L} - t] = B_{\frac{1}{\epsilon}}[\mathcal{L}' - t'] \right\}.$$

Here $B_R[\mathcal{L}] = \mathcal{L} \cap B_R(0)$ is the R-patch of $\mathcal{L}$ at 0, i.e., the patch defined by the closed R-ball at 0.

2.4. Background on lattices, Model sets and Meyer sets

Meyer sets can be thought of as (quite natural) generalizations of lattices. They have been introduced by Meyer in [54] in the purely theoretical context of 'expanding sets via Fourier transforms'. After the discovery of quasicrystals they have become a most prominent class of examples for such structures. Our discussion follows [56, 57, 68] to which we refer for further details and references. For the topic of regular complete Meyer sets we also highlight [58], which gives an introduction into the topic by surveying the results of [6].

A *lattice* in a group G is a uniformly discrete subgroup such that the quotient of G by this subgroup is compact. Thus, any lattice is a Delone set. It is not hard to see that a Delone set L is a lattice if and only if it satisfies

$$L - L = L.$$

Obviously, this gives that $L - L$ is uniformly discrete and hence locally finite. Thus, a lattice is a Delone set of finite local complexity. Moreover, whenever L is a lattice then

$$L^* := \{\chi \in \widehat{G} : (\chi, x) = 1 \text{ for all } x \in L\}$$

[2] A *net* in a topological space X is a function from a directed set A to X: the image of $\alpha \in A$ is denoted x_α and the net is denoted $(x_\alpha)_\alpha$. The net *converges* to $x \in X$ if for each neighborhood U of x there is $\beta \in A$ so that $x_\alpha \in U$ for all $\alpha \geq \beta$. Convergence of nets completely describes the topology of X – see [34].

is a lattice in $\widehat{G}$. It is called the *dual lattice*. Thus, whenever L is a lattice the set of its ε-dual characters

$$L^{\varepsilon} := \{\chi \in \widehat{G} : |(\chi, x) - 1| \leq \varepsilon \text{ for all } x \in L\}$$

is relatively dense. Meyer sets are generalizations of lattices. It turns out that they can be characterized by suitable relaxations of each of the features discussed so far. In fact, each of the features given in the next theorem can be seen as a weakening of a corresponding feature of a lattice.

Theorem 2.2. *Let $\mathcal{L}$ be a Delone set in G. Then, the following assertions are equivalent:*

(i) $\mathcal{L} - \mathcal{L} \subset \mathcal{L} + F$ *for some finite set* $F \subset G$.

(ii) *For any* $\varepsilon > 0$ *the set*

$$\mathcal{L}^{\varepsilon} = \{\chi \in \widehat{G} : |(\chi, x) - 1| \leq \varepsilon \text{ for all } x \in \mathcal{L}\}$$

of ε-dual characters of $\mathcal{L}$ is relatively dense in $\widehat{G}$.

(iii) *There exists a cut-and-project scheme $(H, \tilde{L})$ over G together with a compact W with $\mathcal{L} \subset \curlywedge(W)$.*

Here, the last point requires some explanation. A *cut-and-project scheme over G* consists of a locally compact abelian group H together with a lattice $\widetilde{L}$ in $G \times H$ such that the two natural projections $p_1 \colon G \times H \longrightarrow G$, $(t, h) \mapsto t$, and $p_2 \colon G \times H \longrightarrow H$, $(t, h) \mapsto h$, satisfy the following properties:

- The restriction $p_1|_{\tilde{L}}$ of p_1 to $\tilde{L}$ is injective.
- The image $p_2(\tilde{L})$ is dense in H.

Let $L := p_1(\tilde{L})$ and $(.)^{\star} \colon L \longrightarrow H$ be the mapping $p_2 \circ (p_1|_{\tilde{L}})^{-1}$. Note that * is indeed well defined on L. Given an arbitrary (not necessarily compact) subset $W \subset H$, we define $\curlywedge(W)$ via

$$\curlywedge(W) \ := \ \{x \in L : x^{\star} \in W\}.$$

The set W is then sometimes referred to as the *window*.

The preceding discussion explains all notation needed in the third point of the above theorem. While we do not give a complete proof of the theorem here, we will include some explanation in order to give the reader some of the ideas involved. In particular, we will provide a sketch of how (iii) implies (i). To do so we first highlight two very crucial features of the construction via $\curlywedge$.

Proposition 2.3. *Let $(H, \widetilde{L})$ be a cut-and-project scheme over G. Then,*

- $\curlywedge(W)$ *is relatively dense if the interior of W is non-empty.*
- $\curlywedge(W)$ *is uniformly discrete if the closure of W is compact.*

Now, whenever a Delone set $\mathcal{L}$ is contained in $\curlywedge(W)$ with W compact, then $\mathcal{L} - \mathcal{L}$ is contained in $\curlywedge(W) - \curlywedge(W) \subset \curlywedge(W - W)$. As $W - W$ is compact, we infer uniform discreteness of $\mathcal{L} - \mathcal{L}$. In fact, this argument can be extended to give that any set of the form $\mathcal{L} \pm \cdots \pm \mathcal{L}$ (with finitely many terms) is uniformly discrete.

We can now provide a proof for (iii)$\Longrightarrow$ (i) as follows. As $\mathcal{L}$ is a Delone set, there exists a compact $K \subset G$ with $\mathcal{L} + K = G$. By (iii) and the argument we just gave we have that

$$F := (\mathcal{L} - \mathcal{L} - \mathcal{L}) \cap K$$

is finite. Consider now arbitrary $x, y \in \mathcal{L}$. Then,

$$x - y = z + k$$

for some $k \in K$ and $z \in \mathcal{L}$ (as $\mathcal{L} + K = G$). Now k satisfies $k = x - y - z$ and hence belongs to $\mathcal{L} - \mathcal{L} - \mathcal{L}$ as well. These considerations show that k belongs to the finite set F, finishing the proof.

Remark. If the group G is compactly generated (as is the case for $G = \mathbb{R}^N$) even more is known. In this case the Delone set $\mathcal{L}$ satisfies $\mathcal{L} - \mathcal{L} \subset \mathcal{L} + F$ for some finite F if and only if $\mathcal{L} - \mathcal{L}$ is uniformly discrete. This was first shown by Lagarias for $G = \mathbb{R}^N$ [39]. As discussed in [6] the proof carries over to compactly generated G.

Definition 2.4 (Meyer set). A Delone set $\mathcal{L} \subset G$ is called Meyer if it satisfies one of the equivalent properties of the preceding theorem. The dynamical system induced by a Meyer set is called a Meyer dynamical system.

It is worth noting that any Meyer set is an FLC-Delone set. Indeed, it is a Delone set by definition. Moreover, by the first property in Theorem 2.2 it satisfies $\mathcal{L} - \mathcal{L} \subset \mathcal{L} + F$ for a finite F. This immediately implies that $\mathcal{L} - \mathcal{L}$ is locally finite and, hence, $\mathcal{L}$ has finite local complexity. Let us also emphasize that the Meyer property is substantially stronger than the FLC property.

Meyer sets can be further distinguished depending on properties of W:

An *inter model set*, associated with the cut-and-project scheme $(H, \widetilde{L})$, is a non-empty subset $\mathcal{L}$ of G of the form

$$x + \curlywedge(y + W^\circ) \subset \mathcal{L} \subset x + \curlywedge(y + W),$$

where $x \in G$, $y \in H$, and $W \subset H$ is compact with

$$W = \overline{W^\circ}.$$

Such an inter model set is called *regular* if the Haar measure of the boundary ∂W of W is zero.

Observe that any inter model set is a Meyer set, i.e., a Delone set contained in some $\curlywedge(W)$ for W compact. Indeed, by construction it is contained in some $\curlywedge(W)$ with W compact and the Delone property follows from Proposition 2.3.

Our results below show that it is worth providing special names for inter model sets which are repetitive.

Definition 2.5 (Complete Meyer set). A set $\mathcal{L}$ in G is called a complete Meyer set if it is repetitive and there exists a cut and project scheme $(H, \widetilde{L})$ and a compact $W \subset H$ with $W = \overline{W^\circ}$ such that $x + \curlywedge(y + W^\circ) \subset \mathcal{L} \subset x + \curlywedge(y + W)$. If furthermore the Haar measure of the boundary of W is zero, the set $\mathcal{L}$ is called a regular complete Meyer set.

Remark. Given $y \in H$, if $y + \partial W$ does not intersect L^* one has $\curlywedge(y + W^\circ) = \curlywedge(y + W)$ and so the inter model set defined by such parameters y and $x \in G$ equals $x + \curlywedge(y + W)$ and is repetitive. It is then referred to as a model set (or a cut-and-project set). Any repetitive model set is thus a complete Meyer set. In order to carefully state our results also in the case of singular choices for y (that is, y for which $y + \partial W$ meets L^*) we use the new name 'complete Meyer set'. So strictly speaking, some complete Meyer sets are not model sets (for the given CPS and window)[3]. Furthermore, the name complete Meyer set suggests a process of completing a Meyer set, namely to a repetitive inter model set it sits in. Concrete ideas about that can be found in [2].

Remark. Regular complete Meyer sets have been a prime source of models for quasicrystals. The reason is that these sets have (pure) point diffraction and this is a characteristic feature of quasicrystals. Indeed, a rigorous mathematical framework for diffraction was given by Hof in [28]. In this work it is also shown that regular complete Meyer sets have a lot of point diffraction. That this spectrum is pure was then shown later in [29, 68]. These works actually prove that the dynamical systems arising from regular complete Meyer sets have pure point dynamical spectrum (with continuous eigenfunctions). This is then combined with a result originally due to Dworkin [17] giving that pure point dynamical spectrum implies pure point diffraction. Recent years have seen quite some activity towards a further understanding of this result of Dworkin. In this context we mention results on a converse (i.e., that pure point diffraction implies pure point spectrum) obtained in [46, 5, 26, 53]. The relevance of regular complete Meyer sets in the study of quasicrystals has been underlined by recent results showing that the other main class of examples – those arising from primitive substitutions – is actually a subclass if pure point diffraction is assumed (compare Corollary 4.26 and subsequent discussion).

As is clear from the discussion, any lattice L in G is a Meyer set as well. In this case we can take H to be the trivial group and $W = H$ and this gives that a lattice is a regular complete Meyer set. Later we will also encounter Delone sets $\mathcal{L}$ whose periods

$$P(\mathcal{L}) := \{t \in G : \mathcal{L} - t = \mathcal{L}\}$$

form a lattice. Such a Delone set is called *completely periodic* (or *crystalline*). Note that the periods of a Delone set automatically form a discrete subgroup of G. Thus, the requirement of complete periodicity is really that the periods form a relatively dense set. It is not hard to see that any completely periodic Delone set $\mathcal{L}$ has the form

$$\mathcal{L} = L + F$$

[3]In some works the terminology repetitive model set is used for any element in the hull of a model set with non-singular parameter y. With this usage a complete Meyer set and a repetitive model set are the same thing.

with a lattice L (viz the periods) and a finite set F. This easily gives that any completely periodic Delone set is a Meyer set again.

Any Meyer set over G gives rise to a rotation on an compact abelian group. This will play quite a role in the subsequent analysis. By (iii) of Theorem 2.2, given a Meyer set there exists a cut-and-project scheme $(H, \tilde{L})$ over G. As $\tilde{L}$ is a discrete and co-compact subgroup of $G \times H$, the quotient

$$\mathbb{T} := (G \times H)/\tilde{L}$$

is a compact abelian group and there is a natural group homomorphism

$$G \longrightarrow \mathbb{T},\ t \mapsto (t, 0) + \tilde{L}.$$

In this way, there is natural action of G on $\mathbb{T}$ and $(\mathbb{T}, G)$ is a rotation over G. The system $(\mathbb{T}, G)$ is sometimes referred to as the canonical torus associated to the cut-and-project scheme [8].

3. Equicontinuous factors and Delone dynamical systems

In this section we will deal with special dynamical systems. When defining the concepts of equicontinuity and of almost periodicity we will assume that the topology of X comes from a metric d. We will then refer to the corresponding dynamical systems as *metrizable dynamical systems*. This assumption of metrizability is not necessary for the subsequent results to hold. In fact, in order to formulate the concepts and prove the results, it suffices to have a topology generated by a uniform structure. In particular, the results apply to Delone dynamical systems on arbitary locally compact, σ-compact abelian groups (as these systems can be topologized by a uniform structure [68]). So, in order to simplify the presentation we will define the concepts in the metric case only, but state the results for the general case. We refer the reader to Auslander's book [4] for more detailed information and further background.

3.1. Equicontinuous actions

In this section we recall some of the theory of equicontinuous systems. Further aspects related to proximality will be discussed later in Section 4.

Consider a minimal dynamical system (X, G, α) with X a compact metric space and G a locally compact abelian group acting by α on X. If the action is free (that is, $\alpha_t(x) = x$ for some $x \in X$ implies t is the identity of G) then X can be seen as a compactification of G: it is the completion of one orbit and this orbit is a copy of G. One might ask when is X a group compactification, that is, when does X carry a group structure such that the orbit is a subgroup isomorphic to G, or, in other words, when is (X, G) a rotation on a compact abelian group?

An answer to this question can be given in terms of equicontinuity.

Definition 3.1 (Equicontinuous system). The metrizable dynamical system (X, G, α) is called equicontinuous if the family of homeomorphisms $\{\alpha_t\}_{t\in G}$ is equicontinuous, i.e., if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$d(\alpha_t(x), \alpha_t(y)) < \epsilon$$

for all $t \in G$ and all $x, y \in X$ with $d(x, y) < \delta$.

Remark. The definition of equicontinuity may seem to depend on the particular choice of the metric. However, by compactness of the underlying space X it turns out to hold for one metric if and only if it holds for any metric (which induces the topology).

An equicontinuous system admits an invariant metric which induces the same topology. Indeed one can just take

$$\overline{d}(x, y) := \sup_{t\in G} d(t \cdot x, t \cdot y).$$

Likewise, any compact metrizable abelian group $\mathbb{T}$ admits a left invariant metric: Whenever d is a metric then $\overline{d}(x, y) := \sup_{t\in\mathbb{T}} d(x - t, y - t)$ is a metric on $\mathbb{T}$ which is invariant. This similarity is not a coincidence: According to the following well-known theorem (see, for example, [37]) equicontinuous minimal systems are rotations on groups.

Theorem 3.2 (Equicontinuous systems as rotations on compact groups). *The minimal dynamical system (X, G) is equicontinuous if and only if it is conjugate to a minimal rotation on a compact abelian group.*

Sketch of proof. A rotation on a compact abelian group is obviously equicontinuous. Conversely, if (X, G) is equicontinuous, given any point $x_0 \in X$ the operation $t_1 \cdot x_0 + t_2 \cdot x_0 := (t_1 + t_2) \cdot x_0$ extends to an addition in X so that X becomes a group with x_0 as neutral element. □

Remark. An equicontinuous dynamical system need not be minimal but a transitive equicontinuous dynamical system (i.e., one containing a dense orbit) is always minimal. So in the context of Delone (and tiling – see Section 4.3.4) dynamical systems, which are by definition the closure of one orbit, equicontinuous systems are always minimal.

Equicontinuity is strongly related to almost periodicity. In order to explain this further we will need some notation. Let (X, G) be a metrizable dynamical system and d the metric on X. The ϵ-ball around $x \in X$ is denoted by $B_\epsilon(x)$. The elements of

$$\mathcal{R}(x, \epsilon) := \{t \in G : t \cdot x \in B_\epsilon(x)\}$$

are called *return vectors* to $B_\epsilon(x)$. Now, (X, G) is called *uniformly almost periodic* if, for any $\epsilon > 0$ the joint set of return vectors to ϵ-balls, given by

$$A = \bigcap\nolimits_{x\in X} \mathcal{R}(x, \epsilon)$$

is relatively dense (i.e., there exists a compact K with $A + K = G$).

Theorem 3.3 ([4], Equicontinuity via almost periodicity). *The minimal dynamical system (X, G) is equicontinuous if and only if it is uniformly almost periodic.*

For Delone dynamical systems we can even be more specific. Recall that a continuous bounded function f on G is called *Bohr almost periodic* if for any $\varepsilon > 0$ the set

$$\{t \in G : \|f - f(\cdot - t)\|_\infty \leq \varepsilon\}$$

of its *almost ε-periods* is relatively dense.

Theorem 3.4 (Characterization equicontinuous Delone systems). *Let $\mathcal{L}$ be a Delone set in the locally compact, σ-compact abelian group G. Then, the following assertions are equivalent:*

(i) *The function*

$$f_{\mathcal{L},\varphi} : G \longrightarrow \mathbb{C}, \;\; f_{\mathcal{L},\varphi}(t) = \sum_{x \in \mathcal{L}} \varphi(t - x)$$

is Bohr-almost periodic for any continuous and compactly supported function φ on G.

(ii) *The hull $\Omega_{\mathcal{L}}$ is a compact abelian group with neutral element $\mathcal{L}$ and group addition satisfying $(\mathcal{L} - t) + (\mathcal{L} - s) = (\mathcal{L} - s - t)$ for all $t, s \in G$.*

(iii) *The dynamical system $(\Omega_{\mathcal{L}}, G)$ is equicontinuous.*

The theorem (and a proof) can be found in [32]. Of course, the equivalence between (ii) and (iii) follows from the above Theorem 3.2. The equivalence between (i) and (ii) is close in spirit to Theorem 3.3. However, the proof given in [32] is based on [51].

It is possible to describe (up to conjugacy) all equicontinuous systems over G via subgroups of $\widehat{G}$. The reason is basically that an equicontinuous system is a rotation on a compact group due to Theorem 3.2. This compact group in turn is determined by its dual group, which is just a subgroup of $\widehat{G}$. Moreover, the elements of this dual group turn out to be just the continuous eigenvalues of the system. This highlights the role of the continuous eigenvalues.

As it is both instructive in itself and also enlightening for the material presented in the next section we now give a more detailed discussion of these connections. We start with a general construction and then state the result describing the equicontinuous minimal systems over G.

Let $\mathcal{E}$ be a subgroup of $\widehat{G}$. We equip $\mathcal{E}$ with the discrete topology and denote the dual group of $\mathcal{E}$ by $\mathbb{T}_{\mathcal{E}}$, i.e.,

$$\mathbb{T}_{\mathcal{E}} := \widehat{\mathcal{E}}.$$

Then, $\mathbb{T}_{\mathcal{E}}$ is a compact abelian group. The inclusion $\mathcal{E} \hookrightarrow \widehat{G}$ gives rise (by Pontryagin duality) to a group homomorphism $j : G \longrightarrow \mathbb{T}_{\mathcal{E}}$ with dense range. In this way there is a natural action of G on $\mathbb{T}_{\mathcal{E}}$ via

$$G \times \mathbb{T}_{\mathcal{E}} \longrightarrow \mathbb{T}_{\mathcal{E}}, \;\; (t, x) \mapsto t \cdot x := j(t)x,$$

where on the right-hand side the elements $j(t)$ and x of $\mathbb{T}_{\mathcal{E}}$ are just multiplied via the group multiplication of $\mathbb{T}_{\mathcal{E}}$. This action can be explicitly calculated as

$$(t \cdot x)(\chi) = (\chi, t)(x, \chi)$$

for $\chi \in \mathcal{E}$. As j has dense range, this group action is minimal and hence (as $\mathbb{T}_{\mathcal{E}}$ is compact), it is uniquely ergodic with the (normalized) Haar measure on $\mathbb{T}_{\mathcal{E}}$ as the invariant measure. Thus, $(\mathbb{T}_{\mathcal{E}}, G)$ is a minimal uniquely ergodic rotation on a compact group. By Theorem 3.2, such a system is equicontinuous. Furthermore, countability of $\mathcal{E}$ is equivalent to metrizability of the dual group $\mathbb{T}_{\mathcal{E}}$ by standard harmonic analysis.

Theorem 3.5. *Let G be a locally compact abelian group.*

(a) *Let (X, G) be an equicontinuous minimal dynamical system. Then, (X, G) is conjugate to $(\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}, G)$.*

(b) *Whenever $\mathcal{E}$ is a subgroup of $\widehat{G}$ then $(\mathbb{T}_{\mathcal{E}}, G)$ is the unique (up to conjugacy) equicontinuous minimal dynamical system whose set of continuous eigenvalues is given by $\mathcal{E}$.*

Proof. A sequence of claims establishes the statements of the theorem.

Claim 1. Let $(\mathbb{T}, G)$ be a minimal rotation on the compact group $\mathbb{T}$ with action of G induced by the homomorphism $j : G \longrightarrow \mathbb{T}$. Let $\iota : \widehat{\mathbb{T}} \longrightarrow \widehat{G}$ be the dual of j (i.e., $\iota(\chi)(t) = (\chi, j(t))$ for $\chi \in \widehat{\mathbb{T}}$ and $t \in G$). Then, ι is injective and the set $\mathcal{E}$ of continuous eigenvalues of $\mathbb{T}$ is just the image of $\widehat{\mathbb{T}}$ under ι. In this way, the system $(\mathbb{T}, G)$ is completely determined by the set of its continuous eigenvalues.

Proof. The action of G on $\mathbb{T}$ comes from a homomorphism $j : G \longrightarrow \mathbb{T}$ with dense range. Thus, its dual map $\iota : \widehat{\mathbb{T}} \longrightarrow \widehat{G}$ is injective. We therefore have to show that the continuous eigenvalues of $(\mathbb{T}, G)$ are just given by the $\iota(\widehat{\mathbb{T}})$. This in turn follows easily from the definitions: Any χ in the dual group of $\mathbb{T}$ gives rise to the continuous $f_\chi : \mathbb{T} \longrightarrow \mathbb{C}, f_\chi(x) = (\chi, x)$ which takes the value 1 at the neutral element of $\mathbb{T}$. This f_χ is an eigenfunction to $\iota(\chi)$ as

$$f_\chi(t \cdot x) = (\chi, j(t)x) = (\chi, x)(\chi, j(t)) = (\iota(\chi), t) f_\chi(x).$$

Conversely, whenever λ is a continuous eigenvalue of $(\mathbb{T}, G)$ with associated eigenfunction f_λ we can assume without loss of generality that f_λ takes the value 1 on the neutral element of $\mathbb{T}$. Then the eigenvalue equation gives that f_λ is actually an element of the dual group of $\mathbb{T}$.

Claim 2. Consider the map ι from the previous claim as a bijective group homomorphism from the discrete group $\widehat{\mathbb{T}}$ onto the discrete group $\mathcal{E}$. Let κ be its dual mapping $\mathbb{T}_{\mathcal{E}} = \widehat{\mathcal{E}}$ to $\mathbb{T}$. Then κ establishes a topological conjugacy between $\mathbb{T}_{\mathcal{E}}$ and $\mathbb{T}$.

Proof. This follows directly from unwinding the definitions.

Claim 3. Let $\mathcal{E}$ be a subgroup of $\widehat{G}$ and $(\mathbb{T}_{\mathcal{E}}, G)$ be the associated equicontinous dynamical system. Then the set $\mathcal{E}_{\text{top}}(\mathbb{T}_{\mathcal{E}}, G)$ of continuous eigenvalues of $(\mathbb{T}_{\mathcal{E}}, G)$ can naturally be identified with $\mathcal{E}$.

Proof. The previous claim shows that the set of continuous eigenvalues is just given by the dual of the group $\mathbb{T}_{\mathcal{E}}$. By construction this dual is just $\mathcal{E}$.

The statements of the theorem follow directly from the above claims. □

Remark. Part (b) of the previous theorem concerns the construction of equicontinuous systems with a given group of continuous eigenvalues. In the context of the present paper it seems appropriate to point out that (for subgroups of the Euclidean space) such a system can even be constructed as the torus of a cut-and-project scheme [64].

3.2. Maximal equicontinuous factor

In this section we consider a topological minimal dynamical system (X, G). There exists a largest (in a natural sense) equicontinuous factor of this system. It is known as *maximal equicontinuous factor*. This factor can be obtained in various ways including

- via the dual of the topological eigenvalues,
- via a quotient construction,
- via the Gelfand spectrum of continuous eigenfunctions.

All of this is certainly well known. In fact, substantial parts can be found in the book [4] for example. Other parts of the theory, while still quite elemenary, seem to be scattered over the literature. In particular, in the context of our interests, corresponding constructions are discussed in [6, 11]. For the convenience of the reader we give a rather detailed discussion here.

Recall that $\mathcal{E}_{\text{top}}(X, G)$ carries the discrete topology. Thus we are exactly in the situation discussed at the end of the previous section. In particular, there is a group homomorphism $j : G \longrightarrow \mathbb{T}_{\mathcal{E}_{\text{top}}(X,G)}$ with dense range. This group homomorphism induces an action of G on $\mathbb{T}_{\mathcal{E}_{\text{top}}(X,G)}$ making it into a rotation on a compact group. In this way we obtain a minimal uniquely ergodic dynamical system $(\mathbb{T}_{\mathcal{E}_{\text{top}}(X,G)}, G)$ out of our data. Explicitly, the action of G on $\mathbb{T}_{\mathcal{E}_{\text{top}}(X,G)}$ is given as

$$(t \cdot x)(\chi) = (\chi, t)(x, \chi)$$

for $\chi \in \mathcal{E}_{\text{top}}(X, G)$. We present two remarkable properties of this dynamical system in the next two lemmas.

Lemma 3.6 (Description of equicontinuous factors). *Let (X, G) be a minimal dynamical system and (Y, G) an equicontinuous factor. Then $\mathcal{E}_{\text{top}}(Y, G)$ is a subgroup of $\mathcal{E}_{\text{top}}(X, G)$ and (Y, G) is a factor of $(\mathbb{T}_{\mathcal{E}_{\text{top}}(X,G)}, G)$.*

Proof. Let $\pi : X \longrightarrow Y$ be the factor map. Then any continuous eigenfunction f on Y (to the eigenvalue χ) gives rise to the eigenfunction $f \circ \pi$ on X (to the

eigenvalue χ). This shows the first part of the statement. Obviously, the embedding

$$\mathcal{E}_{\mathrm{top}}(Y,G) \longrightarrow \mathcal{E}_{\mathrm{top}}(X,G)$$

is a group homomorphism. Dualising, we obtain a group homomorphism

$$\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)} = \widehat{\mathcal{E}_{\mathrm{top}}(X,G)} \longrightarrow \widehat{\mathcal{E}_{\mathrm{top}}(Y,G)} = \mathbb{T}_{\mathcal{E}_{\mathrm{top}}(Y,G)}.$$

This group homomorphism can easily be seen to be a G-map. Hence, by compactness and minimality of the groups in question, it is onto and hence a factor map. The desired statement now follows as $\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(Y,G)}$ is conjugate to (Y,G) by Theorem 3.5. □

Remark. Note that the factor map from $(\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}, G)$ to (Y,G) can be chosen to be a group homomorphism. This is clear from the proof. In fact, it is a general phenomenon: As is easily shown, if a rotation on a compact group is a factor of another rotation on a compact group (mapping the neutral element to the neutral element), then the factor map is a group homomorphism (see also the proof of (b) of Theorem 3.1 in [51] for this type of reasoning).

Lemma 3.7. *Let (X,G) be a minimal dynamical system. Then $(\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}, G)$ is a factor of (X,G).*

Proof. Fix an arbitrary point $x_0 \in X$ and choose for any $\chi \in \mathcal{E}_{\mathrm{top}}(X,G)$ a continuous eigenfunction f_χ with $f_\chi(x_0) = 1$. Using the eigenfunction equation along the orbit of x_0 and minimality we have:

$$f_{\chi\eta} = f_\chi f_\eta, \quad \overline{f_\chi} = f_{-\chi}.$$

Any $x \in X$ then gives rise to the map

$$\widehat{x} : \mathcal{E}_{\mathrm{top}}(X,G) \longrightarrow S^1, \widehat{x}(\chi) := f_\chi(x).$$

By construction and the choice of the f_χ the map $\widehat{x}$ is a character on $\mathcal{E}_{\mathrm{top}}(X,G)$, i.e., an element of $\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}$. It is not hard to see that the map

$$X \longrightarrow \mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}, x \mapsto \widehat{x},$$

is a G-map. By minimality and compactness, it is then a factor map. □

The two previous lemmas establish the following theorem.

Theorem 3.8. *Let (X,G) be minimal. Then there exists a unique (up to conjugacy) factor $(X_{\max}, G)$ of (X,G) satisfying the following two properties:*

- *The factor $(X_{\max}, G)$ is equicontinuous.*
- *Whenever (Y,G) is an equicontinuous factor of (X,G) then (Y,G) is a factor of $(X_{\max}, G)$.*

Proof. Existence follows directly from the previous two lemmas. Uniqueness can be shown as follows: Let (Y_1,G) and (Y_2,G) be two factors with the above properties. Then both Y_1 and Y_2 are compact groups and there exist factor maps $\pi_{1,2} : Y_1 \longrightarrow Y_2$, and $\pi_{2,1} : Y_2 \longrightarrow Y_1$. Without loss of generality we can assume that $\pi_{1,2}$ maps

the neutral element e_1 of Y_1 onto the neutral element, e_2, of Y_2 and $\pi_{2,1}$ maps e_2 to e_1 (otherwise we could just compose the maps with appropriate rotations of the groups). Then

$$\pi_{1,2} \circ \pi_{2,1}(e_1) = e_1 \qquad \text{and} \qquad \pi_{2,1} \circ \pi_{1,2}(e_2) = e_2.$$

As both $\pi_{1,2}$ and $\pi_{2,1}$ are G-maps, this shows that $\pi_{1,2} \circ \pi_{2,1}$ agrees with the identity on Y_1 on the whole orbit of e_1 and $\pi_{2,1} \circ \pi_{1,2}$ agrees with the identity on Y_2 on the whole orbit of e_2. By continuity of the maps and denseness of the orbits we infer that the maps are inverse to each other. This shows that indeed (Y_1, G) and (Y_2, G) are conjugate. □

Definition 3.9. The factor $(X_{\max}, G)$ is called the maximal equicontinuous factor of (X, G). The corresponding factor map will be denoted by $\pi_{\max}$.

We now present two additional ways to view the maximal equicontinuous factor.

A construction via quotients. Let the equivalence relation $\sim$ on X be defined by $x \sim y$ if and only if $f(x) = f(y)$ for every continuous eigenfunction f and let

$$\pi : X \longrightarrow X/\sim =: X_\sim$$

be the canonical projection. Let $X_\sim$ have the quotient topology so that a map g on $X_\sim$ is continuous if and only if $g \circ \pi$ is continuous. It is not hard to see that the action of G on X induces an action of G on $X_\sim$ by the (well-defined!) map

$$G \times X_\sim \longrightarrow X_\sim, (t, \pi(x)) \mapsto (\pi(t \cdot x)).$$

Then π is a G-map and hence a factor map. Note that the preceding considerations show that whenever f_χ is a continuous eigenfunction to the eigenvalue χ there exists a unique continuous eigenfunction g_χ on $X_\sim$ with $f_\chi = g_\chi \circ \pi$.

If we are given additionally an invariant probability measure m on X, this measure is transferred to a G-invariant measure $m_\sim := \pi(m)$ on $X_\sim$. In this way we have constructed a dynamical system $(X_\sim, G)$ together with an invariant measure $m_\sim$.

A construction via the Gelfand transform. Let $\mathcal{A}$ be the closed (w.r.t to $\|\cdot\|_\infty$) subalgebra of $C(X)$ generated by the continuous eigenfunctions. Then $\mathcal{A}$ is a commutative C^*-algebra and there exists therefore a compact space $X_\mathcal{A}$ and a continuous isomorphism of algebras (Gelfand transform)

$$\Gamma : \mathcal{A} \longrightarrow C(X_\mathcal{A}).$$

The space $X_\mathcal{A}$ is in fact nothing but the set of all multiplicative linear non-vanishing functionals on $\mathcal{A}$ and the map Γ is then given by

$$\Gamma(f)(\phi) = \phi(f)$$

for $f \in \mathcal{A}$ and $\phi \in X_\mathcal{A}$. The action of G on X induces an action of G on $\mathcal{A}$ and this in turn induces an action of G on $X_\mathcal{A}$. By construction, Γ is then a G-map

with respect to these actions, i.e.,

$$\Gamma(f(t\cdot)) = (\Gamma f)(t\cdot).$$

Assume now that we are additionally given an invariant probability measure m on X. Then m can be seen as a linear positive functional $m : \mathcal{A} \longrightarrow \mathbf{C}, f \mapsto m(f)$. Thus, via Γ, it induces a linear positive functional $m_{\mathcal{A}}$ on $C(X_{\mathcal{A}})$ and hence $m_{\mathcal{A}}$ is a measure on $X_{\mathcal{A}}$. It is not hard to see that the map $\Gamma : \mathcal{A} \longrightarrow C(X_{\mathcal{A}})$ extends to a unitary G-map

$$U : L^2_{pp,\mathrm{top}}(X, m) \longrightarrow L^2(X_{\mathcal{A}}, m_{\mathcal{A}}),$$

where $L^2_{pp,\mathrm{top}}(X, m)$ is the subspace of $L^2(X, m)$ generated by the continuous eigenfunctions. (The subscript *pp* in the notation refers to 'pure point'.) As is easily seen, the only G-invariant functions on $L^2(X_{\mathcal{A}}, m_{\mathcal{A}})$ are constant. Thus $m_{\mathcal{A}}$ is an ergodic measure on $(X_{\mathcal{A}}, G)$ and we have expressed $L^2_{pp,\mathrm{top}}$ as the L^2-space of a dynamical system.

We now discuss how all three constructions give the same dynamical system (up to conjugacy).

Theorem 3.10. *Let (X, G) be a minimal dynamical system. Then the dynamical systems $(X_{\mathcal{A}}, G)$, $(\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}, G)$ and $(X_{\sim}, G)$ are canonically conjugate. In particular, they are all uniquely ergodic and minimal and have pure point spectrum.*

Proof. Chose for any $\chi \in \mathcal{E}_{\mathrm{top}}(X, G)$ a continuous eigenfunction f_χ. Let g_χ on $X_\sim$ be the unique function with $f_\chi = g_\chi \circ \pi$. Fix an $x_0 \in X$ arbitrarily. We can assume without loss of generality that $f_\chi(x_0) = 1$ for all $\chi \in \mathcal{E}_{\mathrm{top}}(X, G)$. Write $\mathbb{T}$ for $\mathbb{T}_{\mathcal{E}_{\mathrm{top}}(X,G)}$.

We first show the statement on canonical conjugacy. To do so we provide explicit maps: Define

$$J : X_\sim \longrightarrow \mathbb{T} \text{ via } J(\pi(x))(\chi) = g_\chi(\pi(x)) = f_\chi(x).$$

Then J indeed maps $X_\sim$ into $\mathbb{T}$ (as we had normalized our f_χ with $f_\chi(x_0) = 1$). Obviously, J is continuous and injective. As it is a G-map and the action of G on $\mathbb{T}$ is minimal, the map J has dense range. As $\mathbb{T}$ is compact, J is then a homeomorphism.

We now turn to proving that $X_\sim$ and $X_{\mathcal{A}}$ are homeomorphic. Consider

$$\Pi : C(X_\sim) \longrightarrow C(X), g \mapsto g \circ \pi.$$

Then $B := \Pi(C(X_\sim))$ is a closed subalgebra of $C(X)$. As $C(X_\sim)$ is generated by the g_χ, $\chi \in \mathcal{E}_{\mathrm{top}}(X, G)$, the algebra B is generated by $f_\chi = g_\chi \circ \pi$, $\chi \in \mathcal{E}_{\mathrm{top}}(X, G)$. Therefore $B = \mathcal{A}$ and thus Π gives an isomorphism between $C(X_\sim)$ and $\mathcal{A}$. Dualising Π, we obtain a homeomorphism between $X_\sim$ and $X_{\mathcal{A}}$. It is easy to check that all maps involved are G-maps.

The last statement of the theorem is clear since $(\mathbb{T}, G)$ is uniquely ergodic, minimal, and has pure point spectrum. □

The preceding considerations are **summarized** as follows: We have given three different constructions of a certain topological factor of (X, G). This factor is given by a rotation, i.e., an action of the group G on a compact group via a group homomorphism with dense range from G to the compact group. The L^2-space of this factor corresponds to the part of the L^2-space of the original dynamical system coming from continuous eigenfunctions.

In general it is not easy to decide whether a given equicontinuous factor is the maximal equicontinuous factor. However, there is one sufficient condition which is of considerable relevance.

Lemma 3.11. *Let (X, G) be a minimal dynamical system and (Y, G) an equicontinuous dynamical system. If (Y, G) is a factor of (X, G) with factor map π and there exists $y \in Y$ such that $\pi^{-1}(y)$ consists of only one point, then (Y, G) is the maximal equicontinuous factor.*

Proof. Let $(X_{\max}, G)$ be the maximal equicontinuous factor with corresponding factor map $\pi_{\max}$. By Theorem 3.8, the dynamical system (Y, G) is then a factor of $(X_{\max}, G)$. Denote the corresponding factor map by π_Y. Without loss of generality we can then assume that π_Y maps the neutral element of $X_{\max}$ to the neutral element of Y and that

$$\pi = \pi_Y \circ \pi_{\max}$$

(otherwise we can just compose π_Y and π with suitable rotations). We will show that π_Y is a homeomorphism. As π_Y is a factor map, it is onto and continuous. It therefore suffices to show that it is one-to-one. So, let $p, q \in Y$ be given with $\pi_Y(p) = \pi_Y(q)$. As discussed in a remark above, the map π_Y is a group homomorphism. Thus, we obtain

$$\pi_Y(gp) = \pi_Y(g)\pi_Y(p) = \pi_Y(g)\pi_Y(q) = \pi_Y(gq)$$

for all $g \in Y$. As π_Y is onto, we can now chose $g \in Y$ with $\pi_Y(gp) = y = \pi_Y(gq)$. As, by the assumption of the lemma,

$$\pi^{-1}(y) = \pi_{\max}^{-1}(\pi_Y^{-1}(y))$$

consists of only one point, we obtain from the last equality that $gp = gq$ and, hence, $p = q$. This is the desired injectivity. □

The class of dynamical systems appearing in the previous lemma is rather important (for us, in describing regularity properties of Meyer sets – see Theorems 3.16 and 3.19 below) and has a name of its own.

Definition 3.12 (Almost-automorphic system). Let (X, G) be a minimal dynamical system and $(X_{\max}, G)$ its maximal equicontinuous factor. If there exists a $y \in X_{\max}$ such that $\pi_{\max}^{-1}(y)$ has only one element then (X, G) is called almost-automorphic.

We finish this section with a discussion of local freeness in our context. This will be relevant in the discussion in Section 4.3. The equivalence of (i) and (ii) in

the following result can be found in [11]. We include a complete proof (as [11] only contains a proof of one direction).

Whenever the group G acts on X an element t of G is said to act *freely* if $t \cdot x \neq x$ for all $x \in X$. The action is called *locally free* if there exists a neighborhood U of $e \in G$ such that any $t \in U \setminus \{e\}$ acts freely. If U can be chosen as G the action is called *free*.

Lemma 3.13. *Let $\mathcal{E}$ be a subgroup of $\widehat{G}$ and consider the associated dynamical system $(\mathbb{T}_{\mathcal{E}}, G)$ with action given by*

$$(t \cdot x)(\chi) = (\chi, t)(x, \chi)$$

for $x \in \mathbb{T}_{\mathcal{E}}$ and $\chi \in \mathcal{E}$. Then the following assertions are equivalent:

(i) *The action is locally free.*
(ii) *The quotient $\widehat{G}/\overline{\mathcal{E}}$ is compact, where $\overline{\mathcal{E}}$ is the closure of $\mathcal{E}$ in $\widehat{G}$.*
(iii) *The stabiliser $\{t \in G : t \cdot x = x \text{ for (some) all } x\}$ of the action is a discrete subgroup of G.*

In particular, the action is free if and only if $\mathcal{E}$ is dense in $\widehat{G}$.

Proof. The equivalence between (i) and (iii) is clear. It remains to show the equivalence between (i) and (ii). By definition $t \in G$ acts freely if and only if $t \cdot x \neq x$ for all $x \in \mathbb{T}_{\mathcal{E}}$, which is the case if and only if $(\chi, t) \neq 1$ for at least one $\chi \in \mathcal{E}$. By continuity, the latter can be rephrased as $(\chi, t) \neq 1$ for at least one $\chi \in \overline{\mathcal{E}}$. Consider now the exact sequence of abelian groups

$$e \to \widehat{\widehat{G}/\overline{\mathcal{E}}} \to G \xrightarrow{q} \hat{\overline{\mathcal{E}}} \to e$$

which is the dual to the exact sequence $e \to \overline{\mathcal{E}} \to \widehat{G} \to \widehat{G}/\overline{\mathcal{E}} \to e$. Let U be an open neighborhood of $e \in G$ and $e \neq t \in U$. Set $\eta = q(t) \in \hat{\overline{\mathcal{E}}}$. Then, $(\eta, \chi) = (\chi, t)$ because q is dual to the inclusion $\mathcal{E} \hookrightarrow \widehat{G}$. Thus, there exists a χ with $(\chi, t) \neq 1$ if and only if η does not belong to $\ker q$. Hence $t \cdot x \neq x$ for all $x \in \mathbb{T}_{\mathcal{E}}$ if and only if $t \notin \ker q$. Thus G acts locally freely if and only if there exists an open $e \in U \subset G$ such that $U \cap \ker q = \{e\}$. By exactness of the sequence above, this the case if and only if $\widehat{\widehat{G}/\overline{\mathcal{E}}}$ is discrete and hence if and only if $\hat{G}/\overline{\mathcal{E}}$ is compact.

Furthermore, G acts freely if and only if the above is true for $U = G$, which is equivalent to $\mathcal{E}$ being dense in $\widehat{G}$. □

3.3. Delone dynamical systems via their maximal equicontinuous factor

In this section we will study dynamical systems arising as the hull of FLC-Delone sets. The basic aim is to characterize features of the Delone set in question by how close its dynamical system is to its maximal equicontinuous factor. A rough description of our results is that the more ordered the set is, the closer its hull is to its maximal equicontinuous factor. More precise statements will be given below. We will distinguish two situations. In one situation we are given a Meyer set and characterize it by features of its hull. In the other situation we are given the hull of an FLC-Delone set.

In order to set the perspective on the results in the next two subsections we briefly recall a 'hierarchy of order' we have encountered within the FLC-Delone sets (see Section 2.4). Let $\mathcal{L}$ be a Delone set. Then $\mathcal{L}$ has finite local complexity if and only if $\mathcal{L} - \mathcal{L}$ is locally finite. $\mathcal{L}$ is a Meyer set if and only if $\mathcal{L} - \mathcal{L} \subset \mathcal{L} + F$ for some finite set $F \subset G$ or, equivalently, if and only if $\mathcal{L}$ is a subset of $x + \curlywedge(W)$ with compact $W \subset H$ and $x \in G$ for some cut-and-project-scheme $(\widetilde{L}, H)$ over G. Now the following classes of Meyer sets $\mathcal{L}$ can be distinguished, each of them defined by a stronger requirement than the previous one:

- $\mathcal{L}$ is a complete Meyer set if it is repetitive with $x + \curlywedge(W^\circ) \subset \mathcal{L} \subset x + \curlywedge(W)$ for some $x \in G$ and some cut-and-project scheme $(\widetilde{L}, H)$ over G and a compact $W \subset H$ with $W = \overline{W^\circ}$.
- $\mathcal{L}$ is a regular complete Meyer set if it is repetitive with $x + \curlywedge(W^\circ) \subset \mathcal{L} \subset x + \curlywedge(W)$ for some $x \in G$ and some cut-and-project scheme $(\widetilde{L}, H)$ over G and a compact $W \subset H$ with $W = \overline{W^\circ}$ and boundary ∂W of Haar measure zero.
- $\mathcal{L}$ is completely periodic if the set $\{t \in G : t + \mathcal{L} = \mathcal{L}\}$ is a lattice or, equivalently, if $\mathcal{L} = L + F$ for a lattice L and a finite set F.

3.3.1. Regularity of Meyer sets via dynamical systems. In this section we study the hull $(\Omega_{\mathcal{L}}, G)$ of a repetitive Meyer set $\mathcal{L}$. We investigate how the hierarchy of Meyer sets discussed above is reflected in injectivity properties of the factor map between this hull and its maximal equicontinuous factor.

Theorem 3.14. *Let G be a locally compact abelian group and $\mathcal{L}$ a repetitive Meyer set in G. The following assertions are equivalent:*

(i) *$\mathcal{L}$ is completely periodic.*
(ii) *The dynamical system $(\Omega_{\mathcal{L}}, G)$ of $\mathcal{L}$ is conjugate to its maximal equicontinuous factor (i.e., each point in the maximal equicontinuous factor has exactly one inverse image point under the factor map).*

Note that (ii) actually says that $(\Omega_{\mathcal{L}}, G)$ is just a rotation on a compact group. Using the material of the last section (on characterizing the maximal equicontinuous factor via a quotient construction) this can be seen to be equivalent to the continuous eigenfunctions separating the points. With this formulation (instead of (ii)) the result is shown in [6]. Of course, the implication (i)$\Longrightarrow$ (ii) is clear and it is the other implication where all the work lies. We will comment a bit on the method of proof after stating the next result.

Theorem 3.15. *Let G be a locally compact abelian group and let $\mathcal{L}$ be a repetitive Meyer set in G. The following are equivalent:*

(i) *$\mathcal{L}$ is a regular complete Meyer set.*
(ii) *The dynamical system $(\Omega_{\mathcal{L}}, G)$ of $\mathcal{L}$ is an almost-1-to-1 extension of its maximal equicontinuous factor (i.e., the set of points in the maximal equicontinuous factor with exactly one inverse image point under the factor map has full measure).*

This is a reformulation of the main result of [6] in terms of the maximal equicontinuous factor. The formulation of (ii) given there is somewhat different and says that the continuous eigenfunctions separate almost all points and the system is uniquely ergodic and minimal. Now, by our discussion on how to obtain the maximal equicontinuous factor (via a quotient construction) and Lemma 3.11, this is just (ii).

The direction (i)$\Longrightarrow$ (ii) was first shown in the special case of the Penrose system by Robinson in [63]. General complete Meyer sets were then treated by Schlottmann in [68]. The main work of [6] is to show the implication (ii)$\Longrightarrow$ (i). There, the *diffraction* of the Meyer set plays a key role. Condition (ii) implies that diffraction is a pure point measure and this can be used to introduce new topologies on the Delone sets. Taking suitable completions of the hull of $\mathcal{L}$ in these topologies one then obtains the ingredients of a cut-and-project scheme via a method of [10]. The main work of [6] is then to prove regularity features of the window. This regularity is shown by an analysis of rotations on compact groups. A crucial role in these considerations is played by continuity of the eigenfunctions. This continuity is related to uniform existence of certain ergodic averages [62, 50]. As such it has also played a major role in the investigation of diffraction and the so-called Bombieri/Taylor conjecture. We refer the reader to [50] for further discussion and background.

Theorem 3.16. *Let $G = \mathbb{R}^N$ and let $\mathcal{L}$ be a repetitive Meyer set in G. The following are equivalent:*

(i) *$\mathcal{L}$ is a complete Meyer set.*
(ii) *The dynamical system $(\Omega_{\mathcal{L}}, G)$ of $\mathcal{L}$ is an almost-automorphic system (i.e., the set of points in the maximal equicontinuous factor with exactly one inverse image point under the factor map is non-empty).*

This is one version of a main result of the Ph.D. thesis of J.B. Aujogue [2]. The implication (i)$\Longrightarrow$ (ii) is somewhat folklore. It is mentioned in the introduction of [64] and can rather directly be derived from Lemma 3.11 and some Baire type arguments (see [68, 6] for related material). This then holds for arbitrary locally compact, σ-compact abelian groups.

The implication (ii)$\Longrightarrow$ (i) is the hard part of the work. It is shown in [2] how to construct a cut-and-project scheme for $\mathcal{L}$ under condition (ii). In fact, the construction of [2] even gives that the associated torus $\mathbb{T} = (G \times H)/\widetilde{L}$ is just the maximal equicontinuous factor whenever the window W satisfies a suitable 'irredundancy' condition.

Under the additional assumptions of unique ergodicity and pure point spectrum, the implication (ii)$\Longrightarrow$(i) can also directly be inferred by combining Theorem 3A and Theorem 6 from [6] (and this then holds in general locally compact σ-compact abelian groups).

3.3.2. Regularity of the hulls of FLC Delone sets. In this section we consider a repetitive FLC Delone set $\mathcal{L}$ in $G = \mathbb{R}^N$ with its dynamical system $(\Omega_{\mathcal{L}}, G)$. It

turns out that the property of this system to be conjugate to a Meyer dynamical system can be characterized via the maximal equicontinuous factor [33]. This characterization provides the crucial additional insight compared to the previous subsection. It allows the derivation of results for FLC Delone sets based on the results for Meyer sets of the last section. As it is (so far) only available for $G = \mathbb{R}^N$ we have to restrict to this situation.

Theorem 3.17 ([33]). *Let $G = \mathbb{R}^N$ and let $\mathcal{L}$ be a repetitive FLC Delone set in G. The following are equivalent:*

(i) *$(\Omega_{\mathcal{L}}, G)$ is conjugate to a Meyer dynamical system.*
(ii) *The system $(\Omega_{\mathcal{L}}, G)$ has at least N linearly independent continuous eigenvalues.*

It is possible to express the result of the previous theorem via the maximal equicontinuous factor.

Corollary 3.18. *Let $G = \mathbb{R}^N$ and let $\mathcal{L}$ be a repetitive FLC Delone set in G. The following are equivalent:*

(i) *$(\Omega_{\mathcal{L}}, G)$ is conjugate to a Meyer dynamical system.*
(ii) *The maximal equicontinuous factor of $(\Omega_{\mathcal{L}}, G)$ has a factor arising from an action of G on $\mathbb{R}^N/\mathbb{Z}^N$ via the mapping $\mathbb{R}^N \longrightarrow \mathbb{R}^N/\mathbb{Z}^N, x \mapsto Ax + \mathbb{Z}^N$ for some invertible linear map $A : \mathbb{R}^N \longrightarrow \mathbb{R}^N$.*
(iii) *The stabilizer of the G-action on the maximal equicontinuous factor of $(\Omega_{\mathcal{L}}, G)$ is a discrete subgroup.*

Proof. The equivalence between (i) and (ii) is a reformulation of Theorem 3.17 based on Theorem 3.5 and the paragraph preceeding it.

It remains to show the equivalence between (i) and (iii). By Lemma 3.13, the condition (iii) is equivalent to compactness of $\widehat{G}/\overline{\mathcal{E}_{\text{top}}}$. Now, for $G = \mathbb{R}^N$, compactness of $\widehat{G}/\overline{\mathcal{E}_{\text{top}}}$ can easily be seen to be equivalent to $\mathcal{E}_{\text{top}}$ containing N linear independent vectors. By the previous theorem this is equivalent to (i). □

Combining the results of the previous section with the preceding theorem yields the following equivalences.

Theorem 3.19 ([2]). *Let $G = \mathbb{R}^N$ and let $\mathcal{L}$ be a repetitive FLC-Delone set in G. The following are equivalent:*

(i) *$(\Omega_{\mathcal{L}}, G)$ is conjugate to a dynamical system arising from a complete Meyer set.*
(ii) *The dynamical system $(\Omega_{\mathcal{L}}, G)$ of $\mathcal{L}$ is an almost-automorphic system, (i.e., the set of points in the maximal equicontinuous factor with exactly one inverse image point under the factor map is non-empty).*

Theorem 3.20 ([33]). *Let $G = \mathbb{R}^N$ and let $\mathcal{L}$ be a repetitive FLC Delone set in G. The following assertions are equivalent:*

(i) *$(\Omega_{\mathcal{L}}, G)$ is conjugate to a dynamical system arising from a regular complete Meyer set.*

(ii) *The dynamical system* $(\Omega_{\mathcal{L}}, G)$ *of* $\mathcal{L}$ *is an almost* 1*-to-*1 *extension of its maximal equicontinuous factor* (*i.e., the set of points in the maximal equicontinuous factor with exactly one inverse image point under the factor map has full measure*).

Formulation of the analogue of the result on complete periodicity in this context leads to the equivalence of the following two assertions for a repetitive FLC Delone set in $G = \mathbb{R}^N$.

- $(\Omega_{\mathcal{L}}, G)$ is conjugate to a Meyer dynamical system arising from a completely periodic set.
- $(\Omega_{\mathcal{L}}, G)$ is conjugate to its maximal equicontinuous factor.

Now, however, it can easily be seen that a dynamical system coming from a repetitive FLC Delone set is conjugate to that of a completely periodic set if and only if the original set is already completely periodic. Also, it is not hard to convince oneself that both systems in question are automatically minimal. Thus there is no need to assume repetitivity. Altogether we then obtain the following result:

Theorem 3.21. *Let* $G = \mathbb{R}^N$ *and let* $\mathcal{L}$ *be a Delone set in* G *with finite local complexity. The following are equivalent:*

(i) *The Delone set* $\mathcal{L}$ *is completely periodic.*
(ii) *The dynamical system* $(\Omega_{\mathcal{L}}, G)$ *agrees with its maximal equicontinous factor.*

Remark. This result had already been proven in [32] using different methods. In fact, the result of [32] is even more general in that it applies to arbitrary compactly generated abelian groups. The result provides an answer to a question of Lagarias [41].

4. Proximality

In the previous section we saw the utility of the maximal equicontiuous factor and its factor map. In the present section we present a different approach to this factor by furnishing an alternative description of the equivalence relation defined by $\pi_{\max}$. Thus, we will start from a dynamical system (X, G) with compact X and abelian G and the factor map $\pi_{\max} : X \longrightarrow X_{\max}$. In order to ease the presentation of certain concepts we will assume metrizability of the system, i.e., of X, in certain places. If the results are valid without this restriction we state them without this restriction (compare also the discussion at the beginning of Section 3). Most of the time the systems will furthermore be required to be minimal.

4.1. Definitions

We consider a variety of relations on X. We begin with the relation induced by $\pi_{\max}$:

1. The *equicontinuous structure relation*

$$R_{\max} := \{(x, y) \in X \times X : \pi_{\max}(x) = \pi_{\max}(y)\}.$$

This relation will be studied by means of the following three relations, each of these is a subset of the equicontinuous structure relation (see the discussion following Theorem 4.2).

2. The *proximality relation* $P := \bigcap_{\epsilon>0} P_\epsilon$ with
$$P_\epsilon := \{(x, y) \in X \times X : \text{there exists } t \in G : d(t \cdot x, t \cdot y) < \epsilon\}.$$
3. The *regional proximality relation* $Q := \bigcap_{\epsilon>0} \overline{P_\epsilon}$, where $\overline{P_\epsilon}$ is the closure of P_ϵ in the product topology.
4. The *syndetic proximality relation* $s^y P$, where we say that x and y are syndetically proximal if for all $\epsilon > 0$,
$$A_\epsilon := \{t \in G : d(t \cdot x, t \cdot y) < \epsilon\}$$
is relatively dense.

The most intuitive of these relations seems to be the proximality relation because it has a direct dynamical meaning: x and y are proximal if they can come arbitrarily close when they are moved around with equal group elements. We cannot expect P to be transitive. Moreover, P need not be a closed relation, i.e., closed in the product topology on $X \times X$. So while P is intuitive, it can be somewhat tricky.

The regional proximality relation looks like an innocent extension of the proximality relation which is guaranteed to be closed. But care has to be taken! While Q contains P it is in general not the smallest closed equivalence relation containing P. It may be non-trivial even if P is the trivial equivalence relation. For later applications it will be useful to spell out that x and y are regional proximal if and only if for any $\varepsilon > 0$ there exist $(x', y') \in P_\varepsilon$ with x' arbitrarily close to x and y' arbitrarily close to y.

The syndetic proximality relation is always transitive, i.e., it is always an equivalence relation. But it is not always closed. It is clearly contained in the proximality relation. Moreover, if the proximality relation is closed then it agrees with the syndetic proximality relation. See [15] for proofs of these facts and more information.

One more thing can be said already about P: if the system is equicontinuous then P must be the trivial relation. The converse is not true, however. A point $x \in X$ is called *distal* if it is only proximal to itself. Systems with trivial proximality relation are called *distal*, because they only have distal points. There exist minimal distal systems which are not equicontinuous, but the celebrated theorem of Furstenberg on the structure of minimal distal systems ([23]) will not concern us here due to the following result. (The result in the stated form is first given in [12] and is generalized to compactly generated groups in [32], where it forms the core of the proof of Theorem 3.21.)

Theorem 4.1 ([12, 32]). *Consider a repetitive non-periodic FLC Delone set in $\mathbb{R}^N$. There exist two distinct elements of its hull which agree on a half-space. In particular the proximality relation on non-periodic FLC Delone systems is non-trivial.*

4.2. Some results for general dynamical systems

We state right away the fundamental result relating the maximal equicontinuous factor and proximality. Again we do not attempt to state it in the most general form.

Theorem 4.2 ([4]). *Let (X, G) be a minimal dynamical system. Then the equicontinuous structure relation $R_{\max}$ is equal to the regional proximal relation Q.*

One direction of containment claimed in this theorem is relatively easy. Note that continuous eigenfunctions of the dynamical system take the same value on proximal pairs: If x is proximal to y and f is a continuous eigenfunction then

$$\frac{f(x)}{f(y)} = \frac{f(t \cdot x)}{f(t \cdot y)} = 1$$

since the first equality is true by the eigenvalue equation for all t, and $t \cdot x$ and $t \cdot y$ can get arbitrarily close when varying t. A similar argument can be employed if x and y are merely regionally proximal – one only needs to take into account an (arbitrarily small) error of ϵ. Thus $Q \subset R_{\max}$. For the other direction one needs to prove that the induced action on X/Q is equicontinuous which is equivalent to showing that there is an invariant metric generating its topology.

Remark. By the previous theorem, the regional proximality relation opens an alternative way to study the topological spectrum. Given that we have a good intution about proximality, this raises the question: How different are the regional proximal and proximal relations? In general they are quite different. For instance, for a topologically weakly mixing system any two elements are regionally proximal [4]. But in the minimal case, x can only be proximal to $t \cdot x$ if they are equal. Indeed, suppose that $\inf_{t' \in \mathbb{R}^N} d(t' \cdot x, t' \cdot t \cdot x) = 0$. Then there exists a sequence $(t_n)_n$ such that $\lim_n d(t_n \cdot x, t_n \cdot (t \cdot x)) = 0$ and so if we take an accumulation point x' of the sequence $(t_n \cdot x)_n$ then $d(x', t \cdot x') = 0$ and so $t \cdot x' = x'$. By minimality we must then also have $t \cdot x = x$.

4.2.1. Distal points. To understand the regional proximal relation it is important to consider points which are regionally proximal but not proximal. Recall that a point $x \in X$ is called distal if it is not proximal to any other point. We denote by X^{distal} the distal points of X. This set might be empty. In fact for metrizable minimal topologically weakly mixing systems it is known [4] (p. 132) that $P^2 = Q = X \times X$, i.e., for any two points $x, y \in X$ there exists a third point $z \in X$ such that x is proximal to z and z is proximal to y. This is only possible if x is not distal. In particular, as we will see that non-Pisot substitution tilings have weakly mixing dynamical systems, their tiling spaces have no distal points. As distality is preserved by the action, X^{distal} is dense if it is not empty. But for metrizable X even more is true: the set X^{distal} is even residual if it is not empty due to a remarkable result of Ellis [20]. Minimal systems for which X^{distal} is not empty are called *point distal*. Veech has extended Furstenberg's structure theorem to point distal systems [75].

Let $\xi \in X_{\max}$. We call $\pi_{\max}^{-1}(\xi)$ the *fiber* of ξ without specifying the map $\pi_{\max}$. A point $\xi \in X_{\max}$ is called *fiber-distal* if all points in its fiber are distal. Since $P \subset R_{\max}$ a point x can only be proximal to a point in the fiber to which it belongs, so ξ is fiber-distal if and only if the proximality relation restricted to its fiber is trivial. We denote the fiber-distal points by $X_{\max}^{\text{distal}}$ and say that a minimal system is fiber-distal if this set is non-empty. Clearly

$$\pi_{\max}^{-1}(X_{\max}^{\text{distal}}) \subset X^{\text{distal}}.$$

Let X^{open} be the set of open points for $\pi_{\max}$, that is, of points x such that $\pi_{\max}$ maps neighbourhoods of x to neighbourhoods of $\pi_{\max}(x)$. X^{open} is always a residual set (and hence non-empty) [75]. A point $\xi \in X_{\max}$ is called fiber-open if all points in its fiber $X_\xi = \pi_{\max}^{-1}(\xi)$ are open. We denote by $X_{\max}^{\text{open}}$ the fiber-open points. We state two fundamental results of Veech [75].

Lemma 4.3 ([75]). *Any distal point is open and so in particular* $X_{\max}^{\text{distal}} \subset X_{\max}^{\text{open}}$. *Furthermore,* $X_{\max}^{\text{open}}$ *is a residual subset of* $X_{\max}$.

In particular $X_{\max}^{\text{open}}$ is always non-empty which shows that the inclusion $X_{\max}^{\text{distal}} \subset X_{\max}^{\text{open}}$ need not be an equality. This is, for instance, the case if the system has no distal points.

Lemma 4.4 ([75]). *If* X^{distal} *is a residual set then there exists a fiber* $\pi_{\max}^{-1}(\xi)$ *with a dense set of distal points.*

4.2.2. Coincidence rank. For our further investigation of the various relations we consider three notions of rank for minimal dynamical systems. The *minimal* and the *maximal rank*, mr and Mr, are the minimal and maximal number of points in a fibre of $\pi_{\max}$, or $+\infty$ if the extrema do not exist. The *coincidence rank*, $\text{cr} \in \mathbb{N} \cup \{+\infty\}$, of a minimal dynamical system counts the maximal number of mutually non-proximal points in a fibre. More precisely, let $\xi \in X_{\max}$ and $\text{card}(\xi, \delta)$ be the maximal number l of elements $x_1, \dots, x_l \in \pi_{\max}^{-1}(\xi)$ such that $(x_i, x_j) \notin P_\delta$, or $\text{card}(\xi, \delta) = +\infty$ if this maximum does not exist. There are a couple of observations to make: First, and this follows from minimality, $\text{card}(\xi, \delta)$ does not depend on ξ, and second, $\text{card}(\xi, \delta)$ is decreasing in δ. So we may define the coincidence rank by

$$\text{cr} = \lim_{\delta \to 0} \text{card}(\xi, \delta) = \sup\{l : \exists x_1, \dots, x_l \in \pi_{\max}^{-1}(\xi), (x_i, x_j) \notin P \text{ for } i \neq j\}$$

and this is independent of $\xi \in X_{\max}$. Here, the first equality is a definition and the second equality follows easily. Moreover, we have the following criterion for finiteness of cr.

Lemma 4.5. *The coincidence rank* cr *is finite if and only if there exists* $\delta_0 > 0$ *such that if* $(x_1, x_2) \in Q \backslash P$ *then* $d(x_1, x_2) \geq \delta_0$.

Indeed, by compactness of the fibers, such a δ_0 cannot exist if cr is infinite, and if cr is finite then the limit $\lim_{\delta \to 0} \text{card}(\xi, \delta)$ must be taken on a strictly positive

value for δ. In particular we see that for finite coincidence rank, $\mathrm{cr} = \mathrm{cr}(\xi, \delta)$ for all ξ and all $0 < \delta \leq \delta_0$.

Independence of $\mathrm{card}(\xi, \delta)$ on ξ also implies that $\mathrm{cr} \leq \mathrm{mr}$ as the coincidence may be measured in a fiber of minimal size. So an almost-automorphic system, i.e., a system with $\mathrm{mr} = 1$, has also $\mathrm{cr} = 1$.

Lemma 4.6. *Consider a metrizable minimal system with finite coincidence rank. If proximality is transitive then for all $\delta > 0$ there exists a compact subset $K \subset G$ such that if $(x, y) \in P$ then $d(t \cdot x, t \cdot y) < \delta$ for some $t \in K$.*

Proof. Let δ_0 be as in Lemma 4.5 and let $0 < \delta \leq \delta_0$. Assume that the statement of the lemma is false, namely that P is transitive but there is a sequence of proximal pairs $(x^{(n)}, y^{(n)}) \in P$ and an ascending chain $(K_n)_n$ of compact subsets $K_n \subset G$ with $\bigcup_n K_n = G$ such that $d(t \cdot x^{(n)}, t \cdot y^{(n)}) \geq \delta$ for all $t \in K_n$. Let $\xi^{(n)} = \pi_{\max}(x^{(n)})$ (which coincides with $\pi_{\max}(y^{(n)})$). Recall that $\mathrm{cr} = \mathrm{card}(\xi, \delta)$ for any ξ and thus $\pi_{\max}^{-1}(\xi^{(n)})$ contains at most cr elements $x_1^{(n)}, \ldots, x_{\mathrm{cr}}^{(n)}$ which are pairwise non-proximal. $x^{(n)}$ must therefore be proximal to one of those, let's say to $x_1^{(n)}$. We may assume that $x_1^{(n)} = x^{(n)}$, as otherwise we could just replace $x_1^{(n)}$ with $x^{(n)}$. By transitivity of P, $y^{(n)}$ cannot be proximal to $x_i^{(n)}$, for $i > 1$. It follows from Lemma 4.5 that $d(t \cdot y^{(n)}, t \cdot x_i^{(n)}) \geq \delta_0 > \delta$ for all t and $i > 1$. We now take sub-sequences of the above sequences such that they all converge, tending to the limits $x^{(\infty)} = x_1^{(\infty)}$, $x_i^{(\infty)}$, for $i = 2, \ldots, \mathrm{cr}$, and $y^{(\infty)}$, respectively. It follows that $d(t \cdot y^{(\infty)}, t \cdot x_i^{(\infty)}) \geq \delta$ for all t and all i, and $d(t \cdot x_i^{(\infty)}, t \cdot x_j^{(\infty)}) \geq \delta$ for all t and all $i \neq j$. Hence all these limit points, which are in the same fiber, are pairwise non-proximal. This fiber contains thus at least $\mathrm{cr} + 1$ distal points. This is a contradiction to the definition of the coincidence rank and so we conclude that there exists a compact $K \subset G$ such that if $(x, y) \in P$ then $d(t \cdot x, t \cdot y) < \delta$ for some $t \in K$. □

The coincidence rank furnishes a criterion for when proximality (P) coincides with regional proximality (Q).

Theorem 4.7 ([11]). *Let (X, G) be a minimal system. Then* $\mathrm{cr} = 1$ *if and only if $P = Q$.*

Corollary 4.8. *Let (X, G) be a minimal system. If (X, G) is almost-automorphic then* $\mathrm{cr} = 1$ *and $P = Q$.*

4.2.3. Consequences of finite coincidence rank. The following consequence of finite coincidence rank will prove to be useful.

Lemma 4.9. *Consider a minimal dynamical system with free $\mathbb{R}^N$-action. If the coincidence rank is finite, the $\mathbb{R}^N$-action on its maximal equicontinuous factor is free.*

Proof. We denote the orbit $\{t\cdot x : t \in \mathbb{R}^N\}$ of x by $Orb_{\mathbb{R}^N} x$. Let H be the stabilizer of a fiber $\pi_{\max}^{-1}(\xi)$ under the $\mathbb{R}^N$-action. Since the action is free $tx \neq sx$ for $s \neq t$ in G and any x in the dynamical system. Thus, the cardinality of H is bounded by the cardinality of $Orb_{\mathbb{R}^N} x \cap \pi_{\max}^{-1}(\xi)$ for any $x \in \pi_{\max}^{-1}(\xi)$. Now in the remark after Theorem 4.2 we have already discussed that x cannot be proximal to a translate of it, as the action is free. Hence all the points in $Orb_{\mathbb{R}^N} x \cap \pi_{\max}^{-1}(\xi)$ are mutually non-proximal. This means that the cardinality of $Orb_{\mathbb{R}^N} x \cap \pi_{\max}^{-1}(\xi)$ is bounded by cr. Thus H is finite. Since $\mathbb{R}^N$ has no finite subgroups except the trivial group, the $\mathbb{R}^N$ action on the maximal equicontinuous factor must be free. □

Note that if cr is finite, then a point $\xi \in X_{\max}$ is fiber distal if and only if its fiber contains exactly cr elements.

The following yields criteria for when the coincidence rank equals the minimal rank. We need metrizability of X as we make use of the result of Ellis [20] mentioned above.

Proposition 4.10. *Let (X, G) be a metrizable minimal system and suppose* $\mathrm{cr} < \infty$. *The following are equivalent:*

(i) $\mathrm{cr} = \mathrm{mr}$,
(ii) $X_{\max}^{\mathrm{distal}} \neq \emptyset$,
(iii) *the system is point distal ($X^{\mathrm{distal}} \neq \emptyset$),*
(iv) $X_{\max}^{\mathrm{distal}} = X_{\max}^{\mathrm{open}}$.

Proof. The equivalence between (i) and (ii) is rather direct, given that each fiber must have at least cr elements which are not mutually proximal. Clearly (ii) implies (iii). We now show that (iii) implies (i): If the system is point distal then, by the already mentioned result of Ellis [20], X^{distal} is residual and hence, by Veech's Lemma 4.4, there exists $\xi \in X_{\max}$ such that $\pi_{\max}^{-1}(\xi)$ contains a dense set of distal points. But if cr is finite a fiber can only have finitely many distal points (namely at most cr). It follows that the fiber $\pi_{\max}^{-1}(\xi)$ is finite and contains only distal points. Thus $\mathrm{mr} \leq \mathrm{cr}$. As the inequality $\mathrm{cr} \leq \mathrm{mr}$ is clear we obtain (i).

Finally, we discuss the equivalence between (ii) and (iv). Here, (ii) follows from (iv) by Veech's Lemma 4.3. It remains therefore to show that if $X_{\max}^{\mathrm{distal}} \neq \emptyset$, then a fiber-open point is fiber-distal. Note that $\pi_{\max}$ being open at x is equivalent to the condition that whenever $(\xi_n)_n$ is a sequence in $X_{\max}$ converging to $\xi := \pi_{\max}(x)$, we can lift that sequence to a sequence $(x_n)_n \subset X$ which converges to x.

So let $X_{\max}^{\mathrm{distal}}$ be non-empty and therefore a dense subset of $X_{\max}$ (it is invariant under the action). Let $\xi \in X_{\max}^{\mathrm{open}}$ and $(\xi_n)_n$ a sequence in $X_{\max}^{\mathrm{distal}}$ converging to ξ. By Lemma 4.5 the points in a fiber of a fiber distal point have mutual distance at least $\delta_0 > 0$. Thus, the set of limits of sequences $(x_n)_n \subset X$ which are convergent and satisfy $\pi_{\max}(x_n) = \xi_n$ can have at most cr points. So by the above criterion for open points $\pi_{\max}^{-1}(\xi)$ cannot contain more than cr points. Thus ξ is fiber distal. □

Theorem 4.11. *Consider a metrizable minimal system (X, G) with finite coincidence rank. The following are equivalent:*

(i) *Proximality is transitive.*
(ii) *Proximality is closed.*
(iii) *Proximality coincides with syndetic proximality.*

If, moreover, G is compactly generated and $X_{\max}$ is connected then the above conditions are equivalent to:

(iv) *Proximality coincides with regional proximality* (*i.e.,* $\mathrm{cr} = 1$).

Remark. The hypothesis that cr is finite is crucial. Indeed, it may happen that P is trivial, and hence closed, whereas Q is not trivial, that is, the system is distal without being equicontinuous. By Theorem 4.7, such systems do not have $\mathrm{cr} = 1$. In fact, by the preceding theorem, such systems must always have infinite coincidence rank.

Proof. (Of Theorem 4.11.) We start with the implication (i) implies (ii). Suppose that there is a sequence of proximal pairs $(x^{(n)}, y^{(n)}) \in P$ tending to $(x^{(\infty)}, y^{(\infty)})$. Let $\delta = \frac{\delta_0}{2}$. By Lemma 4.6 we can find a compact $K \subset G$ such that for each n there exists a $t_n \in K$ for which $d(t_n \cdot x^{(n)}, t_n \cdot y^{(n)}) \leq \delta$. By going over to a subsequence we may assume that $(t_n)_n \subset K$ converges, let's say to $t_\infty \in K$. It follows that $(t_n \cdot x^{(n)}, t_n \cdot y^{(n)})$ is a convergent sequence in Q tending to $(t_\infty \cdot x^{(\infty)}, t_\infty \cdot y^{(\infty)})$. Since Q is closed the latter pair lies also in Q. Moreover, $d(t_\infty \cdot x^{(\infty)}, t_\infty \cdot y^{(\infty)}) \leq \delta < \delta_0$. By Lemma 4.5 this implies that $(x^{(\infty)}, y^{(\infty)}) \in P$ and so P is topologically closed.

The implications (ii) implies (iii) and (iii) implies (i) (syndetical proximality is transitive) are proved in [15].

The equivalence between (ii) and (iv) can be found in [11] – we provide a sketch here. If P is closed then X/P is a compact space which is metrizable. The maximal equicontinuous factor map $\pi_{\max}$ factors therefore as $X \to X/P \xrightarrow{\pi} X/Q = X_{\max}$ and the factor map π is a cr-to-1 map. One can show that π is a local homeomorphism, from which it follows that the system $(X/P, G)$ is equicontinuous. (This latter result needs that G is compactly generated and that $X_{\max}$ is connected.) Hence, by maximality of the equicontinuous factor, $X/P = X_{\max}$. □

Remarks

(a) For $G = \mathbb{R}^N$ the assumptions on G and $X_{\max}$ are satisfied. Indeed, $\mathbb{R}^N$ is compactly generated. Moreover, $\mathbb{R}^N$ is connected and, due to minimality, X, and hence $X_{\max}$, must then be connected as well.
(b) By means of suspension of a $\mathbb{Z}^N$-action to an $\mathbb{R}^N$-action (see, for example, [18]) the result applies with $G = \mathbb{Z}^N$ as well.
(c) One can use Lemma 4.6 to give a direct proof of the implication (i) implies (iii). That lemma states that for each $\delta > 0$ there is a compact set $K \subset G$ such that when $(x, y) \in P$ then $d(t \cdot x, t \cdot y) < \delta$ for some $t \in K$. Since P is an invariant relation this implies that when $(x, y) \in P$ then $A_\delta = \{s \in G : d(s \cdot x, s \cdot y) < \delta\}$ satisfies $A_\delta + K = G$, that is, A_δ is relatively dense. Hence x and y are syndetically proximal.

Theorem 4.12 ([11]). *Consider a minimal system (X, G) with finite coincidence rank* cr*. Suppose that $X_{\max}^{\mathrm{distal}}$ has full Haar measure. Let μ be an ergodic probability measure on X. The following are equivalent:*

(i) $\mathrm{cr} = 1$.
(ii) *The system is an almost* 1*-to-*1 *extension of its maximal equicontinous factor.*
(iii) *The continuous eigenfunctions generate $L^2(X, \mu)$.*

If one of these conditions holds then the system is in fact uniquely ergodic.

Proof. The equivalence between (i) and (ii) is rather obvious. Indeed, (ii) $\Longrightarrow$ (i) is clear (compare Corollary 4.8). The reverse (i)$\Longrightarrow$ (ii) follows directly as, by its very definition, the set $X_{\max}^{\mathrm{distal}}$ consists exactly of those ξ with all points in $\pi_{\max}^{-1}(\xi)$ distal, i.e., those ξ with exactly $\mathrm{cr} = 1$ points in $\pi_{\max}^{-1}(\xi)$. Condition (ii) implies that $\pi_{\max}$ yields an isomorphism between $L^2(X, \mu)$ and $L^2(X_{\max}, \eta)$ which implies (iii) as $C(X_{\max})$ spans $L^2(X_{\max}, \eta)$. To show (iii)$\Longrightarrow$ (i), i.e., that $\pi_{\max}$ can't yield an isomorphism on the level of L^2-spaces in case that $\mathrm{cr} > 1$, we somewhat surprisingly need topology, namely it follows once one has shown that $\pi_{\max}$ is almost everywhere a covering map. □

4.3. Proximality for Delone sets

We now study particular aspects of proximality for Delone sets in Euclidean space $\mathbb{R}^N$. The first result concerns the restrictions imposed by finite coincidence rank on a repetitive Delone set.

Theorem 4.13. *A repetitive FLC Delone set whose dynamical system has finite coincidence rank is topologically conjugate to a Meyer dynamical system. Moreover, if the Delone set is non-periodic* (*no periods*) *then the topological eigenvalues form a dense subgroup of $\hat{\mathbb{R}}^N$.*

Proof. The completely periodic case is trivial. We treat the non-periodic case, leaving the case of fewer than N periods to the reader. In this case the $\mathbb{R}^N$-action on the hull is free and hence, by Lemma 4.9, so also is the $\mathbb{R}^N$-action on the maximal equicontinuous factor. Lemma 3.13 implies now that the eigenvalues are dense and the result follows from Theorem 3.17. □

4.3.1. Examples and open questions for higher coincidence rank. In Section 3.3 we have presented a hierarchy of properties for Delone sets which is based on how large the set of points $\xi \in \Omega_{\max}$ which have unique pre-image under $\pi_{\max}$ is. This characterisation concerns the case of minimal rank $\mathrm{mr} = 1$ (and hence also $\mathrm{cr} = 1$). We now provide examples of Delone sets whose dynamical system has higher coincidence rank.

$1 < \mathrm{cr} < +\infty$. Model sets and periodic sets are now excluded and we know from Theorem 4.13 that the group of eigenvalues is dense. We will see below that any primitive Meyer substitution tiling has finite maximal rank Mr and hence finite coincidence rank. Furthermore, such a substitution tiling system has $\mathrm{cr} = 1$ if and only if its dynamical spectrum is purely discrete (see Theorem 4.23). So examples

falling into the category $1 < \mathrm{cr} < +\infty$ are primitive Meyer substitution tiling systems for which do not have pure point dynamical spectrum. The most famous such example comes from the Thue–Morse substitution $0 \mapsto 01$, $1 \mapsto 10$. Elements of the corresponding substitution tiling space, in which the lengths of both the 0 and 1 tiles agree (using a decoration to distinguish them) produce Meyer sets whose system has coincidence rank $\mathrm{cr} = 2$.

cr $= +\infty$. Clearly, non-trivial Delone systems which are topologically mixing must have infinite coincidence rank. The system of geometrical non-Pisot substitution tilings are of that kind (see below). However there are also examples of Meyer sets whose systems have infinite coincidence rank. For instance the scrambled Fibonacci substitution provides such examples [33]. If we take the tiling version in which the tiles have length $\frac{1+\sqrt{5}}{2}$ and 1 then the system is topologically weakly mixing; there are no topological eigenvalues besides the trivial one. If we give the tiles both length 1 then the tiling is Meyer and the topological eigenvalues form a subgroup of rank 1. But a subgroup of rank 1 cannot be dense and thus the coincidence rank must be infinite.

Open questions. Our findings above suggest the following questions:

1. Does there exists a non-automorphic Delone dynamical system in which the equicontinuous structure relation coincides with proximality? This means that $\mathrm{cr} = 1 < \mathrm{mr}$ and thus the system does not have any distal point.
2. More generally, do there exist Delone dynamical systems with finite coincidence rank which do not have distal points?

4.3.2. Strong proximality and statistical coincidence. For Delone dynamical systems (in $\mathbb{R}^N$) it is possible to introduce stronger versions of proximality and regional proximality based on the so-called combinatorial metric. Although this metric does not generate the topology of the hull, it is quite useful, particularly for Meyer sets. Recall that two Delone sets are close in the combinatorial metric if they agree on a large ball. Accordingly, it is possible to formulate corresponding stronger versions of proximality without reference to the metric but just via agreement on large balls. To simplify this we define, for a Delone set $\mathcal{L}$ and an $R > 0$,

$$B_R[\mathcal{L}] := \mathcal{L} \cap B_R(0).$$

By analogy with the relations 2.,3., and 4. of Section 4.1 we define the following relations:

5. Two Delone sets $\mathcal{L}_1$, $\mathcal{L}_2$ are *strongly proximal* if for all R there exists $t \in \mathbb{R}^N$ such that $B_R[\mathcal{L}_1 - t] = B_R[\mathcal{L}_2 - t]$.
6. Two Delone sets $\mathcal{L}_1$, $\mathcal{L}_2$ are *strongly regional proximal* if for all R there exist $\mathcal{L}_1', \mathcal{L}_2' \in \Omega$, $t \in \mathbb{R}^N$ such that $B_R[\mathcal{L}_1] = B_R[\mathcal{L}_1']$, $B_R[\mathcal{L}_2] = B_R[\mathcal{L}_2']$ and $B_R[\mathcal{L}_1' - t] = B_R[\mathcal{L}_2' - t]$.
7. Two Delone sets $\mathcal{L}_1$, $\mathcal{L}_2$ are *strongly syndetically proximal* if for all R the set of $t \in \mathbb{R}^N$ for which $B_R[\mathcal{L}_1 - t] = B_R[\mathcal{L}_2 - t]$ is relatively dense.

These definitions are indeed strengthenings of the corresponding relations introduced above. Obviously, strong syndetical proximality implies strong proximality which in turn implies strong regional proximality.

We comment that strong proximality for Delone sets is even more intutive than proximality: two Delone sets are strongly proximal if they share arbitrarily large patches. Theorem 4.1 shows that strong proximality is also non-trivial for non-periodic repetitive FLC Delone systems.

We note the following consequence: If $\mathcal{L}_1$ and $\mathcal{L}_2$ are strongly regional proximal elements in the hull Ω of a Delone set $\mathcal{L}$ then $\mathcal{L}_1 - \mathcal{L}_2 \subset (\mathcal{L} - \mathcal{L}) - (\mathcal{L} - \mathcal{L})$. Indeed, pick $x_1 \in \mathcal{L}_1$ and $x_2 \in \mathcal{L}_2$. Choose $R > \max\{|x_1|, |x_2|\}$. Then there are $\mathcal{L}_1', \mathcal{L}_2' \in \Omega$ and $t \in \mathbb{R}^N$ such that $x_1 \in \mathcal{L}_1'$, $x_2 \in \mathcal{L}_2'$ and $t \in \mathcal{L}_1' \cap \mathcal{L}_2'$. Set $v_i = t - x_i$. Then $x_1 - x_2 = v_2 - v_1 \in (\mathcal{L}_2' - \mathcal{L}_2') - (\mathcal{L}_1' - \mathcal{L}_1')$ and the statement follows as $\mathcal{L}' - \mathcal{L}' \subset \mathcal{L} - \mathcal{L}$ for each element $\mathcal{L}'$ in the hull of $\mathcal{L}$.

We will also consider a statistical variant of the above concepts. This requires a notion of density. Recall that a sequence $(\Lambda_n)_{n\in\mathbf{N}}$ of compact sets in G with non-empty interior is called *van Hove* if it exhausts G and if

$$\lim_{n\to\infty} \frac{|\partial^K \Lambda_n|}{|\Lambda_n|} \;=\; 0$$

for every compact K in G, where for $\Sigma \in G$, we set $\partial^K \Sigma := \overline{((\Sigma + K) \setminus \Sigma)} \cup ((\overline{G \setminus \Sigma} - K) \cap \Sigma)$. Existence of such a sequence can be shown for arbitrary locally compact σ-compact abelian G (see, e.g., [68, p. 145]). Given such a sequence Λ, the *upper density* of a subset $B \subset \mathbb{R}^N$ w.r.t. the sequence is given by

$$\overline{\mathrm{dens}}_\Lambda(B) = \limsup_n \frac{\mathrm{vol}(B \cap \Lambda_n)}{\mathrm{vol}(\Lambda_n)}$$

and the lower density by a similar expression in which the lim sup is replaced by the lim inf. A priori, this notion depends on the choice of sequence. If both expressions coincide and are independent of the van Hove sequence, the common value is simply called the density of B. In the sequel we will mostly assume that we have fixed a van Hove sequence and suppress dependence on it in the notation.

For a uniformly discrete set $\mathcal{D}$ in G, we modify the above to define its upper density:

$$\overline{\mathrm{dens}}_\Lambda(\mathcal{D}) = \limsup_n \frac{\sharp(\mathcal{D} \cap \Lambda_n)}{\mathrm{vol}(\Lambda_n)},$$

and for its lower density, replace lim sup by lim inf. Which formulas apply when we speak of density will be clear from the context.

After these preparations we can now introduce the following relation:

8. Two Delone sets $\mathcal{L}_1$, $\mathcal{L}_2$ are *statistically coincident* if the (upper) density of the symmetric difference $\mathcal{L}_1 \triangle \mathcal{L}_2 = (\mathcal{L}_1 \backslash \mathcal{L}_1 \cap \mathcal{L}_2) \cup (\mathcal{L}_2 \backslash \mathcal{L}_1 \cap \mathcal{L}_2)$ vanishes. We denote this relation by SC.

   Statistical coincidence implies strong proximality.

Lemma 4.14. *Suppose that $\mathcal{L}_1$ and $\mathcal{L}_2$ are statistically coincident Delone sets. Then they are strongly proximal.*

Proof. If $\mathcal{L}_1$ and $\mathcal{L}_2$ are not strongly proximal then there exists $R > 0$ such that for all $t \in \mathbb{R}^N$ we have $B_R[\mathcal{L}_1 - t] \neq B_R[\mathcal{L}_2 - t]$. Hence for all t the symmetric difference $B_R[\mathcal{L}_1 - t] \triangle B_R[\mathcal{L}_2 - t]$ contains at least one point. It follows that the lower density of $\mathcal{L}_1 \triangle \mathcal{L}_2$ is bounded from below by $1/\mathrm{vol}(B_R(0))$. □

From this lemma and the inclusions of the relations discussed above we immediately have:

Corollary 4.15. *If the statistical coincidence relation SC coincides with the equicontinuous structure relation for a repetitive Delone dynamical system then all relations* 1.–7. *agree with SC and* $\mathrm{cr} = 1$.

Let now $(\Omega_\mathcal{L}, G)$ be the dynamical system arising from the hull of a Delone set in $\mathbb{R}^N$. Then Lemma 4.14 says that $\pi_{\max}$ can be factored as $\Omega_\mathcal{L} \to \Omega_\mathcal{L}/SC \to \Omega_\mathcal{L}/Q = \Omega_{\max}$. What is the quotient $\Omega_\mathcal{L}/SC$? A priori, we do not even know whether SC is a closed relation on $\Omega_\mathcal{L}$. To investigate this question we consider the so-called autocorrelation hull, following [6].

For fixed $r > 0$, the mixed autocorrelation pseudometric on the space $\mathcal{U}_r$ of all uniformly r-discrete subsets of $\mathbb{R}^N$ is given by

$$d_{SC}(\mathcal{L}_1, \mathcal{L}_2) = \inf\{\epsilon > 0 : \exists t_1, t_2 \in B(0,\epsilon) : \overline{\mathrm{dens}}((\mathcal{L}_1 - t_1) \triangle (\mathcal{L}_2 - t_2)) \leq \epsilon\}.$$

This induces a complete metric on the quotient $\mathcal{U}_r/SC$ which we also denote by d_{SC}. Define $\beta : (\mathcal{U}_r, d) \to (\mathcal{U}_r/SC, d_{SC})$ by $\beta(\mathcal{L}) = [\mathcal{L}]_{SC}$ but mind that the topology on the quotient is, a priori, not the quotient topology. Hence, a priori, β is not continuous. To be more clear about this, we write the restriction of β to the hull $\Omega_\mathcal{L}$ of $\mathcal{L}$ as a composition $\beta\,|_{\Omega_\mathcal{L}} = i \circ \mathrm{id}_{SC} \circ q$,

$$(\Omega_\mathcal{L}, d) \xrightarrow{q} (\Omega_\mathcal{L}/SC, \mathcal{Q}) \xrightarrow{\mathrm{id}_{SC}} (\Omega_\mathcal{L}/SC, d_{SC}) \xhookrightarrow{i} (\mathcal{U}_r/SC, d_{SC})$$

where $\mathcal{Q}$ stands for the quotient topology. So we see that $\beta\,|_{\Omega_\mathcal{L}}$ is continuous (which is known to be the case, for example, for a regular complete Meyer set – see Theorem 9 of [6]) if and only if id_{SC} is a homeomorphism; that is, the quotient topology coincides with the metric topology from d_{SC}.

The closure of the orbit of $[\mathcal{L}]_{SC}$ in $\mathcal{U}_r/SC$ is called the *autocorrelation hull* of the Delone set $\mathcal{L}$ and is denoted by $\mathbb{A}_\mathcal{L}$. This is not necessarily a compact space and even for repetitive $\mathcal{L}$ we need to keep track of the dependence on $\mathcal{L}$, as a locally isomorphic Delone set may, a priori, yield a different autocorrelation hull.

Theorem 4.16. *Consider a repetitive Delone set $\mathcal{L}$. If $\beta\,|_{\Omega_\mathcal{L}}$ is continuous, then $(\mathbb{A}_\mathcal{L}, \mathbb{R}^N)$ is isomorphic to $(\Omega_{\max}, \mathbb{R}^N)$ and $SC = Q$.*

Proof. If $\beta\,|_{\Omega_\mathcal{L}}$ is continuous then $\beta(\Omega_\mathcal{L}) = \mathbb{A}_\mathcal{L}$ and so $\mathbb{A}_\mathcal{L}$ equals $\Omega_\mathcal{L}/SC$ and is compact. It follows that there is a factor map $\mathbb{A}_\mathcal{L} = \Omega_\mathcal{L}/SC \to \Omega_\mathcal{L}/Q = \Omega_{\max}$. But d_{SC} is invariant under translation and therefore $(\mathbb{A}_\mathcal{L}, \mathbb{R}^N)$ is equicontinuous. Thus the above factor map must be the identity and $SC = Q$. □

Thus, by Corollary 4.15, continuous $\beta|_{\Omega_\mathcal{L}}$ implies cr = 1. Under slightly stronger assumptions, namely that $\mathcal{L}$ is a Meyer set and the associated dynamical system is uniquely ergodic, the first statement of the corollary can also directly be inferred from Theorem 7 of [6] (and the description of the maximal equicontinuous factor via continuous eigenfunctions). That theorem then even implies that all eigenvalues are topological and the dynamical spectrum is pure point [6].

4.3.3. Strong proximality and the Meyer property. In this section we apply the theory developed above to Meyer sets.

Lemma 4.17 ([11]). *For repetitive Meyer sets, the strong versions of proximality, regional proximality and syndetic proximality agree with the usual ones.*

As a consequence, repetitive Meyer sets enjoy a stronger form of finite local complexity:

Corollary 4.18. *Consider a repetitive Meyer set and let $R > 0$. Up to translation, there are only finitely many pairs of R-patches $(B_R[\mathcal{L}_1], B_R[\mathcal{L}_2])$ with $\pi_{\max}(\mathcal{L}_1) = \pi_{\max}(\mathcal{L}_2)$.*

Proof. By finite local complexity, there are only finitely possibilities, up to translation, for $B_R[\mathcal{L}_1]$. So the question is: How many relative positions in a pair $(B_R[\mathcal{L}_1], B_R[\mathcal{L}_2])$ may we have? Let $(x_1, x_2) \in (B_R[\mathcal{L}_1], B_R[\mathcal{L}_2])$. Since $\mathcal{L}_1$ and $\mathcal{L}_2$ are strongly regional proximal we have $x_1 - x_2 \in (\mathcal{L} - \mathcal{L}) - (\mathcal{L} - \mathcal{L})$. By the Meyer property the latter set is uniformly discrete; since also $|x_1 - x_2| \leq 2R$, we see that we have only finitely many possibilities for $x_1 - x_2$. □

The notion of coincidence rank becomes more intuitive for repetitive Meyer sets. In fact, if cr $< \infty$ we can combine Lemma 4.5 and Lemma 4.17 to obtain that there exists $R_0 > 0$ such that

$$\mathrm{cr} = \max\{l | \exists \mathcal{L}_1, \ldots, \mathcal{L}_l \in \pi_{\max}^{-1}(\xi) : \forall t \in \mathbb{R}^N, B_R[\mathcal{L}_i - t] \neq B_R[\mathcal{L}_j - t] \text{ for } i \neq j\}$$

where $\xi \in \Omega_{\max}$ and $R \geq R_0$ are arbitrary. This has two interpretations. A priori, one could expect that the maximum on the r.h.s. becomes larger if R gets larger, because disagreement on larger patches is a weaker condition than on smaller patches. But this is not the case as soon as R is larger than a certain threshold value R_0. A second interpretation is that the distinct Meyer sets in the fiber of a fiber distal ξ have at each point $t \in \mathbb{R}^N$ distinct R_0-patches; that is, they are *non-coincident* on R_0-patches. If cr is not finite then the max on the r.h.s. becomes indeed arbitrarily large as R tends to infinity.

Let $n^R(\xi)$ be the number of different R-patches at 0 which occur in the elements of the fibre of ξ, i.e.,

$$n^R(\xi) = \#\{B_R[\mathcal{L}] : \mathcal{L} \in \pi_{\max}^{-1}(\xi)\}.$$

For Meyer sets, this number is finite by (the proof of) Corollary. 4.18 and we derive from the preceding formula for cr that cr $\leq n^R(\xi)$, provided cr is finite and R is large enough ($R \geq R_0$). Note that $\lim_{R\to\infty} n^R(\xi)$ is the cardinality of $\pi_{\max}^{-1}(\xi)$.

The following lemma was stated and proved in [11] under the assumption of finite maximal rank. This assumption can in fact be dropped.

Lemma 4.19. *For repetitive Meyer sets, n^R is upper semi-continuous; that is, the sets $\{\xi : n^R(\xi) \geq k\}$ are closed in $\Omega_{\max}$ for any $k \geq 0$.*

Proof. Note that if $\mathcal{L}_n \to \mathcal{L}$ in $\Omega_{\mathcal{L}}$ then there exists a sequence $(t_n)_n \subset \mathbb{R}^N$, $t_n \to 0$, such that $B_R[\mathcal{L}_n - t_n] = B_R[\mathcal{L}]$ for all sufficiently large n. Now fix $k \in \mathbb{N}$ and suppose that $(\xi_n)_n$ is a sequence in $\Omega_{\max}$ with $n^R(\xi_n) \geq k$ for all n and with $\xi_n \to \xi$. It follows from Corollary 4.18 that there is $\delta > 0$ with the property that if $\mathcal{L}, \mathcal{L}' \in \Omega_{\mathcal{L}}$ are such that $\pi_{\max}(\mathcal{L}) = \pi_{\max}(\mathcal{L}')$ and $B_R[\mathcal{L}] \neq B_R[\mathcal{L}']$, then $B_R[\mathcal{L} - t] \neq B_R[\mathcal{L}' - t'])$ holds true even for all $t, t' \in B_\delta(0)$. Thus, for each n, we may choose $\mathcal{L}_n^1, \ldots, \mathcal{L}_n^k \in \pi_{\max}^{-1}(\xi_n)$ with $B_R[\mathcal{L}_n^i - t] \neq B_R[\mathcal{L}_n^j - t'])$ for $i \neq j$ and $t, t' \in B_\delta(0)$. By passing to a subsequence, we may assume that $\mathcal{L}_n^i \to \mathcal{L}^i \in \pi_{\max}^{-1}(\xi)$ for $i = 1, \ldots, k$. Then $B_R[\mathcal{L}^i] \neq B_R[\mathcal{L}^j]$ for $i \neq j$, so $n^R(\xi) \geq k$ and the set $\{\xi : n^R(\xi) \geq k\}$ is closed. □

Theorem 4.20. *Consider a repetitive Meyer set $\mathcal{L}$ whose dynamical system $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ has finite coincidence rank. If the fiber distal points have full Haar measure then, for any $R \geq R_0$ and any $\xi \in \Omega_{\max}$, there is a subset $A \subset \mathbb{R}^N$ of density 1 such that*

$$\#\{B_R[\mathcal{L} - t] : \pi_{\max}(\mathcal{L}) = \xi\} = \mathrm{cr} \quad \textit{for all } t \in A.$$

Remark. We can summarise the theorem as saying that locally (that is, by inspection of finite patches) and with probability 1 all fibers have cr elements. Consider for instance the system associated with the Thue–Morse substitution, which has coincidence rank 2 (see the chapter on the Pisot Substitution Conjecture in this volume). The maximal equicontinuous factor has one orbit whose fibres have 4 elements. But only near the 'branching locus' can one find 4 different R-patches, otherwise there are only 2.

Proof. Let $D^R = \{\xi \in \Omega_{\max} : n^R(\xi) = \mathrm{cr}\}$. Then $\Omega_{\max}^{\mathrm{distal}} = \bigcap_{R \geq R_0} D^R$ and so the hypothesis implies that $\eta(D^R) = 1$. Here, η denotes the Haar measure on the maximal equicontinuous factor. Now let $\tilde{n}^R = n^R - \mathrm{cr}\, 1_{D^R}$; that is, $\tilde{n}^R$ is 0 on D^R and otherwise the same as n^R. By the preceding lemma, n^R is upper semi-continuous. Hence, $\tilde{n}^R$ is a finite positive linear combination of indicator functions on compact sets. As the maximal equicontinuous factor is uniquely ergodic, we obtain then a uniform inequality in the ergodic theorem (see, e.g., Lemma 4 in [50]). More specifically, we have for all ξ and all $R \geq R_0$

$$\limsup_n \frac{1}{\mathrm{vol}(\Lambda_n)} \int_{\Lambda_n} \tilde{n}^R(t \cdot \xi) dt \leq \int \tilde{n}^R(\xi') d\eta(\xi'),$$

where $\Lambda = (\Lambda_n)_n$ is a van Hove sequence for $\mathbb{R}^N$. Due to $\eta(D^R) = 1$ the right-hand side in the previous inequality is 0. Let $B = \{t \in \mathbb{R}^N : n^R(t \cdot \xi) \neq \mathrm{cr}\}$. Then

$1_B(t) \leq \tilde{n}^R(t \cdot \xi)$ and thus

$$\overline{\mathrm{dens}}_\Lambda(B) = \limsup_n \frac{1}{\mathrm{vol}(\Lambda_n)} \int_{\Lambda_n} 1_B(t)dt \leq \limsup_n \frac{1}{\mathrm{vol}(\Lambda_n)} \int_{\Lambda_n} \tilde{n}^R(t \cdot \xi)dt = 0.$$

Hence the density of B is 0 and $A = B^c$, the complement of B, has the required property. □

Corollary 4.21. *Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two elements in the hull of a regular complete Meyer set. Then $\pi_{\max}(\mathcal{L}_1) = \pi_{\max}(\mathcal{L}_2)$ if and only if they are statistically coincident.*

Proof. By assumption, the hypothesis of the last theorem is satisfied with $\mathrm{cr} = 1$. In particular Q agrees with the strong proximality relation and so one direction follows immediately from Lemma 4.14. It remains to show that if $\mathcal{L}_1$ and $\mathcal{L}_2$ belong to the same fibre of $\pi_{\max}$ then they are statistically coincident. But the last theorem just says that in this case the density of points where $\mathcal{L}_1$ and $\mathcal{L}_2$ agree on an R-ball is 1, or, in other words, the density of points where $\mathcal{L}_1$ and $\mathcal{L}_2$ disagree is 0. Hence $\mathcal{L}_1$ and $\mathcal{L}_2$ are statistically coincident. □

Remark. The statement of the above corollary is quite at the heart of the considerations of [6]. In fact, as mentioned above, Theorem 9 of [6] gives that, for regular complete Meyer sets, the map $\beta|_{\Omega_\mathcal{L}}$ is indeed continuous. When combined with Theorem 7 of [6], we directly obtain the statement of the corollary. This approach actually shows that the result is valid not only in $\mathbb{R}^N$ but for general locally compact σ-compact abelian groups.

4.3.4. Meyer substitutions. So far we have formulated all our results for Delone sets rather than tilings. This does not really make a difference, as the two are related by mutually local derivations. In particular, all the concepts and results translate into the formalism of tilings. In this section we will use the formalism of tilings, because we find it much more convenient and intuitive for substitutions. For us, a *tile* in $\mathbb{R}^N$ is a subset homeomorphic to a compact N-ball and a *tiling* of $\mathbb{R}^N$ is a collection of tiles with disjoint interiors which covers $\mathbb{R}^N$. The set of tiles of a tiling T which intersect non-trivially a compact subset $K \subset \mathbb{R}^N$ is called a *patch*. In particular, the R-patch at $0 \in \mathbb{R}^N$ is the set of tiles which touch the closed R-ball $B_R(0)$; we denote it by $B_R[T]$. The *support* of a patch P, $\mathrm{supp}(P)$, is the set of points covered by the tiles of P.[4]

For a tiling substitution we suppose we have a finite set $\mathcal{A} = \{\rho_1, \ldots, \rho_k\}$ of translationally inequivalent tiles (called *prototiles*) in $\mathbb{R}^N$ and an expanding linear map Λ. A *substitution* on $\mathcal{A}$ with expansion Λ is a function $\Phi : \mathcal{A} \to \{P : P \text{ is a patch in } \mathbb{R}^N\}$ with the properties: for each $i \in \{1, \ldots, k\}$, every tile in $\Phi(\rho_i)$ is a translate of an element of $\mathcal{A}$; and $\mathrm{supp}(\Phi(\rho_i)) = \Lambda(\mathrm{supp}(\{\rho_i\}))$. Such a substitution

[4] One could include the possibility of decorating the tiles in case one wants to distinguish translationally congruent tiles and then distinguish the support of a tile (the points covered by the tile) from the decorated tile.

naturally extends to patches and even tilings whose elements are translates of the prototiles and it satisfies $\Phi(P-t) = \Phi(P) - \Lambda(t)$.

A patch P is *allowed* for Φ if there is an $m \geq 1$, an $i \in \{1, \ldots, k\}$, and a $v \in \mathbb{R}^N$, with $P \subset \Phi^m(\rho_i) - v$. The *substitution tiling space* associated with Φ is the collection Ω_Φ of all tilings T of $\mathbb{R}^N$ such that every finite patch in T is allowed for Φ. Ω_Φ is not empty and, since translation preserves allowed patches, $\mathbb{R}^N$ acts on it by translation. To define a metric on Ω_Φ, we can borrow the metric we've used for Delone sets: Pick a point y_i in the interior of each prototile ρ_i and, for $T \in \Omega_\Phi$, let $\mathcal{L}(T) = \{y_i + x : x + \rho_i \in T\}$. Then set $d(T, T') := d(\mathcal{L}(T), \mathcal{L}(T'))$. (The set $\mathcal{L}(T)$ is called a *set of punctures* of T.)

The substitution Φ is *primitive* if for each pair $\{\rho_i, \rho_j\}$ of prototiles there is a $k \in \mathbb{N}$ so that a translate of ρ_i occurs in $\Phi^k(\rho_j)$. If the translation action on Ω_Φ is free, which is equivalent to saying that each of its elements is a non-periodic tiling, then Φ is said to be *non-periodic*. If all tilings from Ω_Φ are FLC then Φ is said to be FLC. If Φ is primitive, FLC and non-periodic then Ω_Φ is compact in the metric described above, $\Phi : \Omega_\Phi \to \Omega_\Phi$ is a homeomorphism, and the translation action on Ω_Φ is minimal and uniquely ergodic ([1], [70], [72]). In particular, $\Omega_\Phi = \Omega_T := \overline{\{T - v : v \in \mathbb{R}^N\}}$ for any $T \in \Omega_\Phi$. It will be with respect to the unique ergodic measure μ on Ω_Φ when we speak about the dynamical spectrum and L^2-eigenfunctions. In the context of eigenfunctions (non-periodic) substitutions have a rather special feature: All measurable eigenfunctions are continuous. Thus, all eigenvalues are automatically continuous eigenvalues. For symbolic dynamics this result is due to Host [30]. The case at hand is treated by Solomyak [73].

A Meyer substitution is a substitution Φ such that the elements of Ω_Φ are Meyer tilings, that is, they are MLD to a Meyer set. To check this it suffices to check that, for $T \in \Omega_\Phi$, the set of punctures $\mathcal{L}(T)$ is a Meyer set.

It is easily verified that Φ preserves the regional proximality relation and therefore induces a homeomorphism $\Phi_{\max}$ on the maximal equicontinuous factor $\Omega_{\max}$. In particular $\Phi_{\max}$ satisfies a similar equation

$$\Phi_{\max}(\xi - t) = \Phi_{\max}(\xi) - \Lambda(t)$$

from which one concludes, as Λ has no root of unity eigenvalues, that $\Phi_{\max}$ is ergodic w.r.t. Haar measure.

Proposition 4.22 ([11]). *Consider a Meyer substitution tiling system with primitive non-periodic substitution. Then the following hold:*

(a) *The maximal rank is finite.*
(b) $\Omega_{\max}^{fiber}$ *has full measure.*
(c) *Syndetic proximality is a closed equivalence relation.*
(d) *Two distinct tilings of a fiber distal fiber do not share a common tile.*

Proof. We indicate the idea of some of the proofs. Since Λ is expanding there is a $c > 1$ and an $n \in \mathbb{N}$ such that $B_{cR}(0) \subset \Lambda^n(B_R(0))$. Replacing Φ by Φ^n, we may

suppose that $n = 1$. Hence

$$\begin{aligned} n^{cR}(\Phi_{\max}(\xi)) &= \#\{B_{cR}[\Phi(T)] : \pi_{\max}(T) = \xi\} \\ &\leq \#\{\Phi(B_R[T]) : \pi_{\max}(T) = \xi\} \leq n^R(\xi). \end{aligned}$$

From this we see that the maximal rank is bounded by $\sup_{\xi\in\Omega_{\max}} n^0(\xi)$ which is finite by Corollary 4.18. The argument for the fourth statement is based on a similar reasoning.

Since $n^R \leq n^{cR}$ the above shows also that the sets $D^R(m) := \{\xi \in \Omega_{\max} : n^R(\xi) \leq m\}$, $m \in \mathbb{N}$, are invariant under $\Phi_{\max}$. By Lemma 4.19 $D^R(m)$ is open. By ergodicity of $\Phi_{\max}$ therefore, it has measure 1, provided it is not empty. Since $\Omega_{\max}^{\text{distal}} = \bigcap_{R\geq R_0} D^R(\text{cr})$ the second statement follows if we show that $D^R(\text{cr}) \neq \emptyset$.

Consider a fiber which has minimal rank, i.e., $\xi \in \Omega_{\max}$ such that $\pi_{\max}^{-1}(\xi) = \{T_1, \cdots, T_{\text{mr}}\}$. Suppose that for all $r > 0$ there exists $w \in \mathbb{R}^N$ such that $\forall t \in B_r(w)$ we have $n^R(\xi - t) \geq \text{mr}$; that is, all $B_R[T_i - t]$, $1 \leq i \leq \text{mr}$, are distinct. Then we can find two sequences $(r_k)_k \to \infty$ and $(w_k)_k \in \mathbb{R}^N$ such that $(T_i - w_k)_k$ converge in Ω, let's say to S_i, and $(\xi - w_k)_k$ converges in $\Omega_{\max}$, to ζ, say, and $\forall t \in B_{r_k}(0)$ all $B_R[T_i - w_k - t]$, $1 \leq i \leq \text{mr}$, are distinct. Taking $k \to \infty$ we conclude that all $B_R[S_i - t]$, $1 \leq i \leq \text{mr}$, $t \in \mathbb{R}^N$, are distinct. In particular, the S_i belong to the fiber of ζ and are pairwise non-proximal and so $\text{cr} \geq \text{mr}$. This shows that $\text{cr} = \text{mr}$ and hence $D^R(\text{cr})$ is not empty.

It remains to argue that our assumption is satisfied. So let us suppose the contrary, namely that there exists $r > 0$ such that for all $w \in \mathbb{R}^N$ there exists $t \in B_r(w)$ with $n^R(\xi - t) \leq \text{mr} - 1$. It follows that the lower density of points $t \in R^n$ such that $n^R(\xi - t) \leq \text{mr} - 1$ is strictly positive. Since for all t we have that $n^R(\xi - t) \leq \text{mr}$ (ξ lies in a fiber of rank mr) the ergodic theorem implies that $\int_{\Omega_{\max}} n^R(\xi) d\eta(\xi) < \text{mr}$. Hence $D^R(\text{mr} - 1)$ can't have measure 0. So it must have measure 1. But then $\bigcap_{R\geq R_0} D^R(\text{mr} - 1)$ has measure 1 and so there must be a fiber of rank at most $\text{mr} - 1$ which contradicts the minimality of mr. This shows the second statement.

For the third statement see [11]. □

As a consequence of the previous proposition and our earlier results (in particular, Theorem 4.12 and Corollary 4.15) we have the following list of equivalent characterizations of pure point (measure) dynamical spectrum which hold for primitive non-periodic Meyer substitutions.

Theorem 4.23. *Consider the system of a primitive non-periodic Meyer substitution. The following are equivalent:*

(i) *The (measure) dynamical spectrum is purely discrete.*
(ii) *The dynamical system is an almost 1-to-1 extension of its maximal equicontinuous factor.*
(iii) *The coincidence rank* cr *is* 1.
(iv) *The (strong) proximality relation is closed.*

(v) *The (strong) proximality relation coincides with the equicontinuous structure relation* $R_{\max}$.
(vi) *The (strong) proximality relation coincides with the (strong) syndetic proximality relation.*

We finally present a result of Lee and Solomyak which, for a particular class of substitution tilings, characterizes those which are Meyer substitutions. This class is defined by some further conditions[5] on the linear expansion Λ, namely that

1. Λ is diagonalizable (over $\mathbb{C}$),
2. all eigenvalues of Λ are algebraically conjugate,
3. all eigenvalues of Λ have all the same multiplicity.

It should be said that the eigenvalues of the linear expansion are algebraic integers ([35],[48]) and, if the expansion is diagonalizable they form a union of families ([36]). Here, a family is the set $F_{p,c}$ of roots of a monic, irreducible, integer polynomial p which have absolute value greater than some real number $c > 0$. In other words, a family is a subset of the set of algebraic conjugates of some algebraic integer which can be characterized by the property that if it contains λ then it contains all conjugates which have absolute value at least as large as that of λ. The family $F_{p,c}$ is called a *Pisot family* if $c = 1$ and p has no roots of absolute value 1.

Theorem 4.24 ([47]). *Consider a primitive non-periodic N-dimensional FLC substitution with expansion Λ satisfying the above three properties. The following are equivalent:*

(i) *The substitution is Meyer*
(ii) *The eigenvalues of Λ form a Pisot family.*
(iii) *The continuous eigenvalues of the $\mathbb{R}^N$-action on the hull are dense in $\hat{\mathbb{R}}^N$.*
(iv) *The maximal equicontinuous factor is non-trivial.*

Recall from the discussion at the beginning of Section 4.2.1 that triviality of the maximal equicontinuous factor implies absence of distal points. Given this, we can combine the previous theorem with Proposition 4.22 and Theorem 4.13 to obtain the following strong dichotomy.

Corollary 4.25. *Consider a primitive non-periodic FLC substitution with expansion Λ satisfying the above three properties. Either the system has no distal points, or the distal points form a set of full measure. In the first case the dynamical point spectrum is trivial and in the second the continuous eigenvalues are dense.*

As a consequence of the above discussion we also obtain the following remarkable statement: When it comes to the question of which point sets or tilings have pure point spectrum, all examples produced by substitutions could also be obtained by the cut-and-project formalism. More specifically, the following holds.

[5] It has been announced in [38] that these conditions can be considerably weakened: see the discussion of Pisot families in the chapter on the Pisot Substitution Conjecture in this volume for more detail.

Corollary 4.26. *Suppose that Φ is a primitive non-periodic N-dimensional FLC substitution with expansion Λ satisfying the above three properties. Let $T \in \Omega_\Phi$ and let $\mathcal{L}(T)$ be a set of punctures. Then $(\Omega_\Phi, \mathbb{R}^N)$ has pure point spectrum if and only if $\mathcal{L}(T)$ is a regular complete Meyer set.*

Proof. Let $(\Omega_\Phi, \mathbb{R}^N)$ have pure point spectrum. As all eigenvalues are continuous, we infer that the maximal equicontinuous factor is non-trivial. By Theorem 4.24 the substitution must then be Meyer. Hence, Theorem 4.23 gives that the dynamical system is an almost $1:1$ extension of its maximal equicontinuous factor. Now, Theorem 3.20 implies that $\mathcal{L}(T)$ is a regular complete Meyer set.

As for the converse direction, we note that any complete regular Meyer set gives rise to a dynamical system which is an almost $1:1$ extension of it's maximal equicontinuous factor by Theorem 3.20. From Theorem 4.23 we then infer pure point spectrum. □

Remark. Of course, it is well known that a regular complete Meyer set gives rise to a dynamical system with pure point spectrum (see, e.g., discussion in Section 2.4). The main part of the corollary is thus the converse implication. It has been shown for one-dimensional systems by Sing [69]. For higher-dimensional self-similar substitutions is has been obtained by Lee in [45]. Note, however, that the work of Lee does not seem to claim regularity of the Meyer set but just its completeness. See also the chapter on the Pisot Substitution Conjecture in this volume.

5. Ellis semigroup

If the action of a group G on a space X is transitive we can view X and its maximal equicontinuous factor $X_{\max}$ as two distinct compactifications of the acting group G, the difference arising from the topology in which it is compactified. In this section we consider a third compactification of G – the Ellis semigroup $E(X,G)$ of the dynamical system (X,G). As a space and dynamical system it tends to be by far the most complicated of the three compactifications. But it has one advantage; namely, it naturally carries the structure of a monoid (i.e., a semigroup with neutral element). It therefore offers the possibility to characterize dynamical systems by means of this algebraic structure. There are only a few types of systems for which this has been successfully carried out; non-periodic Delone systems are, however, not among these. So as a first step we simply present some explicit examples of Ellis semigroups coming from Delone sets and observe that they exhibit a very interesting algebraic structure, which we have not seen before in this context. The examples we present are associated with almost canonical cut-and-project patterns. Almost canonical cut-and-project patterns are complete Meyer sets whose windows are polyhedral satisfying further conditions. Surprisingly, the Ellis semigroup for such dynamical systems has a very particular form. It is a completely regular semigroup (or a union of groups) [60]. Furthermore, its

idempotents from a submonoid which is reminiscent of the so-called face semigroup associated to a hyperplane arrangement [14].

The material of Section 5.1 is mostly based on the book of Auslander [4], although most of it can also be found in the book of Ellis [19], and that of the later sections in the thesis of the first author [2, 3].

5.1. Definitions and known properties

An action α of a group G on the compact space X is nothing else than a representation of the group in terms of transformations of X; i.e., for each $t \in G$, α_t is a function from X to X. The set of all functions from X to X is the product set X^X and becomes a compact space when equipped with the Tychonoff topology.

Definition 5.1. The Ellis semigroup $E(X, G)$ is the closure of $\{\alpha_t | t \in G\}$ in X^X.

By definition of the Tychonoff topology, a net $(f_\lambda)_\lambda$ of functions $f_\lambda : X \to X$ converges to some function $f : X \to X$ if for all $x \in X$ the net $(f_\lambda(x))_\lambda$ converges to $f(x)$. Any element of $E(X, G)$ is thus a limit of a net $(\alpha_{t_\lambda})_\lambda$ where $(t_\lambda)_\lambda$ is a net in G. Assuming that the action is faithful we may identify G with $\{\alpha_t | t \in G\}$ and see that $E(X, G)$ is indeed a compactification of G. Furthermore, G acts on E from the left: $\alpha_t^E(f) = \alpha_t \circ f$. Thus $(E(X, G), G)$ is a dynamical system. A factor map $\pi : (X, G) \to (Y, G)$ induces a continuous surjective monoid morphism $\pi_* : E(X, G) \to E(Y, G)$. In fact, the latter is given by the equality $\pi_*(f)(\pi(x)) = \pi(f(x))$ [4][Thm. 7, p. 54].

The basic idea is now to characterize the dynamical system (X, G) by means of the properties of $E = E(X, G)$. We may ask the following questions:

- E consists of functions $f : X \to X$. What are their properties?
- What is the algebraic structure of E?
- E is a compact Hausdorff space. What more can we say about its topology?

Let us elaborate.

E consists of functions $f : X \to X$ which are limits of homeomorphisms. Are these functions still homeomorphisms? If not, are they at least continuous or invertible? We provide an elementary argument why this cannot be the case for all elements in the semigroup of the dynamical system of a repetitive non-periodic FLC Delone set. Recall from Theorem 4.1 that there are two distinct Delone sets $\mathcal{L}^+, \mathcal{L}^-$ in the hull which agree on a half-space. We choose coordinates such that the first component corresponds to the normal into that half-space. Then, whenever $(t_\Lambda)_\lambda$ is a net such that the first component of t_λ tends to $+\infty$, $\mathcal{L}^+ - t_\lambda$ and $\mathcal{L}^- - t_\lambda$ agree on larger and larger balls. By repetitivity we may assume that there are two such nets for which $\lim_\lambda \mathcal{L}^\pm - t_\lambda^\pm = \mathcal{L}^\pm$. By compactness of E we may assume that $\alpha_{t_\lambda^-}$ converges to an element $f \in E$. Then $f(\mathcal{L}^+) = \lim_\lambda(\mathcal{L}^+ - t_\lambda^-) = \mathcal{L}^-$, because eventually the sets $\mathcal{L}^+ - t_\lambda^-$ and $\mathcal{L}^- - t_\lambda^-$ come close. On the other hand $f(\mathcal{L}^+ - t_\lambda^+) = f(\mathcal{L}^+) - t_\lambda^+ = \mathcal{L}^- - t_\lambda^+$ which tends to $\mathcal{L}^+$. Hence if f were continuous the last argument would give $f(\mathcal{L}^+) = \mathcal{L}^+$, a contradiction. So f is not continuous. Furthermore, $f(\mathcal{L}^-) = \lim_\lambda \mathcal{L}^- - t_\lambda^- = \mathcal{L}^-$ and hence f is not injective.

The set E is closed under composition of functions. Accordingly, composition of functions makes E into a semigroup. Moreover, the identity α_0 is a unit for the composition law and so E is a monoid. Care has to be taken, however, concerning the topological properties of the semigroup product. If $(f_\lambda(x))_\lambda$ converges to $f(x)$ then $\lim_\lambda f_\lambda \circ g = f \circ g$ and so the semigroup product is continuous in the left variable: stated differently, right translation by g, $f \mapsto fg := f \circ g$ is a continuous map. It is however, in general, not true that the semigroup product is continuous in the right variable. Moreover, even in case that G is abelian (which is the only case that will concern us) the semigroup product in $E(X, G)$ is, in general, not commutative (although any element of $E(X, G)$ commutes with elements coming from G, i.e., elements of the type α_t).

From the point of view of topology, the Ellis semigroup is either well behaved in the sense that it is separable and every element is the limit of a sequence (as opposed to net) $(\alpha_{t_n})_n$, $t_n \in G$, or it is rather wild in that it contains a homeomorphic image of the Stone–Čech compactification of $\mathbb{N}$ [25]. In the first case the semigroup is called *tame*. If the semigroup is first countable, i.e., has a countable neighborhood base, then it is tame.

One of the important results of the general theory concerns minimal left ideals and the existence of idempotents, that is, elements p satisfying $p^2 = p$. Let F be a closed G-invariant subset of $E = E(X, G)$. Then $\alpha_t^E(F) \subset F$ implies that F is a left ideal of E. It follows also that any minimal left ideal is a minimal component of the dynamical system and hence closed. Furthermore, an application of Zorn's lemma yields that every (closed) minimal left ideal of E contains idempotents.

There are several connections between proximality and the Ellis semigroup. Two that are fundamental are highlighted in the following theorem.

Theorem 5.2. *Consider a compact minimal dynamical system* (X, G).

(a) $x \in X$ *and* $y \in X$ *are proximal if and only if* $p(x) = p(y)$ *for some minimal idempotent of* $E(X, G)$ (*a minimal idempotent is an idempotent in a minimal left ideal*) [4, Thm. 13, p. 89].

(b) *Proximality is transitive if and only if* E *contains a unique minimal left ideal* [4, Thm. 10, p. 88].

Let us note that part (b) of the previous theorem gives uniqueness of the minimal ideal whenever $P = R_{\max}$. In particular, uniqueness of the mimimal ideal holds for minimal systems if $\mathrm{cr} = 1$ (see the discussion in Section 4).

The somewhat nicest case is that in which the system (X, G) is equicontinuous.

Theorem 5.3 ([4, Thm. 3 & 5, p. 52, 53]). *Consider a dynamical system* (X, G). *The system is equicontinuous if and only if* $E(X, G)$ *is a group and all its elements are homeomorphisms. If, moreover, the action is minimal then* $(E(X, G), G)$ *is topologically conjugate to* $(X_{\max}, G)$ *and the conjugacy is a group isomorphism.*

Much more can be said to support the statement: If the semigroup E is well behaved then the system is close to being equicontinuous. For instance, if, for a minimal system, all elements of E are continuous, then they are even homeomorphisms and the system is equicontinuous [21]. On the other hand, if all elements of E are bijective and so E is a group, then the system must be distal and, conversely, triviality of the proximal relation implies that E is a group. Finally we mention that (again for minimal systems) the topology of E is metrizable if and only if the system is equicontinuous [24]. For a non-periodic FLC Delone system, however, the Ellis semigroup is neither a group nor is it metrizable.

Let M be a minimal left ideal and J the set of its idempotents. Then for every $p \in J$ and $m \in M$ we have $mp = m$ and, furthermore, the restriction of the semigroup product to pM makes pM a group with neutral element p [4][Lemma 1, p. 83]. Moreover, all the groups pM are isomorphic, the isomorphism between pM and qM being given by $pm \mapsto qm$; and M is their disjoint union: $M = \bigcup_{p\in J} pM$. Let $\mathcal{G} := p_0 M$ for some chosen $p_0 \in J$. Then, as a semigroup,

$$M \cong \mathcal{G} \times J$$

where we take the product operation on the r.h.s. The semigroup isomorphism is given by $pm \mapsto (p_0 m, p)$. Indeed, J is a sub-semigroup with product given by the so-to-say left domination rule

$$pq = p, \quad p, q \in J$$

and $pmp'm' = pmm'$, showing that the above map preserves the semigroup product. We can say a little bit more about the group $\mathcal{G}$. Since $p_0 E p_0 \subset p_0 M$ and $(\pi_{\max})_*(p_0) = \mathrm{id}$ (p_0 is an idempotent and id is neutral element in $E(X_{\max}, G)$), $(\pi_{\max})_*$ restricts to a surjective group homomorphism $(\pi_{\max})_* : \mathcal{G} \to E(X_{\max}, G)$ and if (X, G) is minimal then the latter is isomorphic to $X_{\max}$.

Lemma 5.4. *Let (X, G) be a minimal dynamical system. Then $(\pi_{\max})_* : \mathcal{G} \to E(X_{\max}, G) \cong X_{\max}$ is an isomorphism if and only if* $\mathrm{cr} = 1$ *($P = R_{\max}$).*

Proof. We first show that $\mathrm{cr} = 1$ implies that $(\pi_{\max})_*$ is an isomorphism. As $(\pi_{\max})_*$ is onto, we only have to show its injectivity. The map $(\pi_{\max})_*$ is injective if $(\pi_{\max})_*(f) = \mathrm{id}$ implies that $f = p_0$. Now $(\pi_{\max})_*(f) = \mathrm{id}$ means that for all $\xi \in X_{\max}$ and $x \in \pi_{\max}^{-1}(\xi)$ we have $f(x) \in \pi_{\max}^{-1}(\xi)$. Let $f \in p_0 M$ be given such that $\pi_{\max}(x) = \pi_{\max}(f(x))$. By the hypothesis $\mathrm{cr} = 1$, the elements x and $f(x)$ are then proximal. Moreover, by (b) of Theorem 5.3, we have that the minimal ideal is unique. From part (a) of that theorem, we then infer that there exists a $p \in J$, such that $pf(x) = p(x)$. Applying p_0 on both sides and using that $p_0 p = p_0$ by the mentioned left domination, and that $f = p_0 f$ by $f \in p_0 M$, we then obtain

$$f(x) = p_0 f(x) = p_0 p f(x) = p_0 p(x) = p_0(x).$$

As x is arbitrary, this shows $f = p_0$.

To prove the converse suppose that $\mathrm{cr} > 1$ so that there are $\xi \in X_{\max}$ and $x, x' \in \pi_{\max}^{-1}(\xi)$ which are not proximal. We may even assume that $p_0(x) = x$ as

we can replace x by p_0x and x' by p_0x' and this will not change non-proximality. (If p_0x and p_0x' were proximal, there would exist, by Theorem 5.3, a $q \in J$ with $qp_0x = qp_0x'$ and this would give $qx = qx'$ and proximality of x and x' would follow from that theorem.) By minimality of the original system and the definition of the Ellis semigroup, there exists $f \in E(X, G)$ such that $f(x') = x$. So if $m = p_0 f$ we have $m(x') = x$. Thus $(\pi_{\max})_*(m)(\xi) = \xi$. Since $E(X_{\max}, G)$ is a group acting fixed point freely on $X_{\max}$, the latter implies that $(\pi_{\max})_*(m) = 0$. But m cannot be an idempotent, because x' is not proximal to x. □

Note that although $(\pi_{\max})_*$ is continuous as a map from $E(X, G)$ to $E(X_{\max}, G)$, one cannot conclude in the above lemma that $\mathcal{G}$ is homeomorphic to $X_{\max}$, as p_0M need not be closed.

In the next section we introduce a family of Delone sets – the almost canonical cut-and-project sets – whose dynamical systems have Ellis semigroups with a particularly nice algebraic description: the entire semigroup, not just the (unique) minimal left ideal, is isomorphic with a sub-semigroup of the product of the maximal equicontinuous factor and a finite monoid of idempotents.

5.2. Almost canonical cut-and-project sets

Almost canonical cut-and-project sets are special types of complete Meyer sets. Their internal group is a real vector space $\mathbb{R}^{N^\perp}$ where $N^\perp$ is called the *codimension* of the set. They are characterized by the form of the set, S, of singular points in the maximal equicontinuous factor $\Omega_{\max} = \mathbb{T} = (\mathbb{R}^N \times \mathbb{R}^{N^\perp})/\tilde{L}$; that is, the set of points $\xi \in \mathbb{T}$ which have a fiber $\pi_{\max}^{-1}(\xi)$ containing more than one point. Recall that S is determined by the boundary points ∂W of the window W, namely

$$S = ((\mathbb{R}^N \times \partial W) + \tilde{L})/\tilde{L} = (\mathbb{R}^N \times (\partial W + p_2(\tilde{L})))/\tilde{L}.$$

Definition 5.5. A complete Meyer set is almost canonical if its internal group H is a vector space $\mathbb{R}^{N^\perp}$ and its window W a finite union of polyhedra and the following two conditions are satisfied:

1. There are finitely many affine hyperplanes $A_i \subset \mathbb{R}^{N^\perp}$, $i \in I$, such that the set $\partial W + p_2(\tilde{L})$ may be alternatively described as
$$\partial W + p_2(\tilde{L}) = \bigcup_{i \in I} A_i + p_2(\tilde{L}).$$
2. Any compact polyhedron whose boundary lies in $\partial W + p_2(\tilde{L})$ can be constructed via a finite sequence of unions, intersections and complements of polyhedra of the form $W + p_2(x)$ for $x \in \tilde{L}$.

Such a situation arises if W is a union of polyhedra whose codimension 1 faces span affine hyperplanes which have a dense stabilizer under the action of $p_2(\tilde{L})$. Then we may take for the A_i these hyperplanes. A so-called canonical cut-and-project set is one for which $W = p_2(C)$ is the projection of a unit cube C for $\tilde{L}$ (w.r.t. to some choice of base for $\tilde{L}$). It satisfies the above criteria since the faces of the projected cube are spanned by lattice vectors.

The advantage of the alternative description of the singular points is that it leads to a very explicit description of the hull $\Omega_{\mathcal{L}}$. In fact

$$\Omega_{\mathcal{L}} = (\mathbb{R}^N \times \mathbb{R}_c^{N^\perp})/\tilde{L}$$

where $\mathbb{R}_c^{N^\perp}$ is a locally compact totally disconnected space, a "cut-up version" of $\mathbb{R}^{N^\perp}$, which is a certain completion of $\mathbb{R}^{N^\perp} \backslash (\bigcup_{i\in I} A_i + p_2(\tilde{L}))$. This allows the calculation of the cohomology groups (see the chapter "Spaces of projection method patterns and their cohomology" in this volume) and of the complexity exponents of the sets [31] and, as we review here, of the Ellis semigroup.

We will not directly look at the Ellis semigroup of the (so-called continuous) dynamical system $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ but first at the semigroup of a reduction of it and obtain $E(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ by suspension. The reduction is obtained from a choice of $N^\perp$-dimensional subspace $F \subset \mathbb{R}^N \times \mathbb{R}^{N^\perp}$ whose intersection with $\tilde{L}$ is a rank $N^\perp$ subgroup $D = F \cap \tilde{L}$. Then F/D is an $N^\perp$ torus in $\mathbb{T}$ which is transversal to the $\mathbb{R}^N$-action. The first return to F/D of the equicontinuous $\mathbb{R}^N$-action on $\mathbb{T}$ yields an equicontinuous $\tilde{L}/D$-action on F/D. By construction $\tilde{L}/D$ is free of rank N and so we simply write it as $\mathbb{Z}^N$. Now let $\Xi = \pi_{\max}^{-1}(F/D)$. This is then transversal to the $\mathbb{R}^N$-action on $\Omega_{\mathcal{L}}$ and the first return yields an action of $\mathbb{Z}^N$ on Ξ. It is the Ellis semigroup of this reduction $(\Xi, \mathbb{Z}^N)$ of $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$ which we now describe more precisely.

We denote now $\Gamma = p_2(\tilde{L})$ and $\Delta = p_2(D)$. Then

$$\Xi \cong \mathbb{R}_c^{N^\perp}/\Delta$$

with $\mathbb{Z}^N$-action induced by Γ, i.e., $\mathbb{Z}^N = \Gamma/\Delta$. Its maximal equicontinuous factor is $\Xi_{\max} = \mathbb{T}^\perp := \mathbb{R}^{N^\perp}/\Delta$.

For each affine hyperplane A_i there is a vector $a_i \in \mathbb{R}^{N^\perp}$ and a codimension 1 subspace $H_i^0 \subset \mathbb{R}^{N^\perp}$ such that $A_i = H_i^0 + a_i$. We then must have that $\bigcap_{i\in I} H_i^0 = \{0\}$. So the dynamical system $(\Xi, \mathbb{Z}^N)$ is entirely described by the data $(\{A_i\}_{i\in I}, \Gamma, \Delta)$ consisting of a finite collection of affine hyperplanes $\{A_i\}_{i\in I}$ in a real vector space $\mathbb{R}^{N^\perp}$ such that the intersection of their corresponding hyperspaces is trivial; a dense rank $N + N^\perp$ sublattice Γ; and a rank $N^\perp$ sublattice Δ which spans $\mathbb{R}^{N^\perp}$. All that follows depends only on this data and does no longer refer to a Delone set or a tiling. The subspace H_i^0 cuts $\mathbb{R}^{N^\perp}$ into two halfspaces. Choose for each i a positive side. We denote by H_i^+ and H_i^- the open half-space on the positive and the negative side, respectively. It is convenient to set $H_i^\infty = \mathbb{R}^{N^\perp}$.

Definition 5.6. The *cut type* $I(h)$ of a point $h \in \mathbb{R}^{N^\perp}$ is the subset

$$I(h) = \{i \in I : h \in A_i + \Gamma\}.$$

A *point type* is a function $\mathfrak{p} \in \{+, -, \infty\}^I$ such that the cone

$$C_{\mathfrak{p}} := \bigcap_{i\in I} H_i^{\mathfrak{p}(i)}$$

is non-empty. We denote by $\mathfrak{P}$ the finite set of point types. We call $C_{\mathfrak{p}}$ the (point) cone associated with $\mathfrak{p}$. Its domain is

$$\operatorname{dom}\mathfrak{p} = \{i \in I : \mathfrak{p}(i) \neq \infty\}.$$

Note that the cone $C_{\mathfrak{p}}$ is a connected component of $\mathbb{R}^{N^\perp} \setminus \bigcup_{i\in\operatorname{dom}\mathfrak{p}} H_i^0$. We denote by $\underline{\infty} \in \{+,-,\infty\}^I$ the function which is constant equal to ∞. Its domain is empty and $C_{\underline{\infty}} = \mathbb{R}^{N^\perp}$. By construction $I(h+\gamma) = I(h)$ for all $\gamma \in \Gamma$ and so the cut type is also defined for the points of the torus $\mathbb{T}^\perp = \mathbb{R}^{N^\perp}/\Delta$.

Recall that the tangent cone $T_S(x)$ at x of a subset $S \subset \mathbb{R}^{N^\perp}$ is the set of vectors v which can be obtained as limits of the form $v = \lim_n \frac{x_n - x}{s_n}$ where $(x_n)_n \subset S$ and $(s_n)_n \subset \mathbb{R}^+$ are sequences such that $\lim_n s_n = 0$. If x lies in the interior of S then $T_S(x) = \mathbb{R}^{N^\perp}$. If C is a closed cone whose tip is at 0 then $T_C(0) = C$.

Theorem 5.7. *The dynamical system $(\Xi, \mathbb{Z}^N)$ is isomorphic to*

$$\Xi = \{(\xi, \mathfrak{p}) \in \mathbb{T}^\perp \times \mathfrak{P} : \operatorname{dom}\mathfrak{p} = I(\xi)\}$$

with $\mathbb{Z}^N$ action given by $t \cdot (\xi, \mathfrak{p}) = (\xi + t, \mathfrak{p})$ and with topology described in terms of convergence of sequences as follows: A sequence $((\xi_n, \mathfrak{p}_n)_n$ converges to $(\xi, \mathfrak{p})$ if and only if $\xi_n \to \xi$ in $\mathbb{T}^\perp$ and eventually $\xi_n - \xi \in \overline{C_{\mathfrak{p}}}$ and $T_{\overline{C_{\mathfrak{p}_n}}}(0) \subset T_{\overline{C_{\mathfrak{p}}}}(\xi_n - \xi)$.

The maximal equicontinuous factor of Ξ is $\mathbb{T}^\perp$ and the factor map is the projection onto the first factor.

The expressions $\xi_n - \xi \subset \overline{C_{\mathfrak{p}}}$ and $T_{\overline{C_{\mathfrak{p}_n}}}(0) \subset T_{\overline{C_{\mathfrak{p}}}}(\xi_n - \xi)$ should be understood for large enough n so that we can lift $\xi_n - \xi$ into a small neighbourhood of 0 in $\mathbb{R}^{N^\perp}$ where the expressions make sense. The condition of inclusion $T_{\overline{C_{\mathfrak{p}_n}}}(0) \subset T_{\overline{C_{\mathfrak{p}}}}(\xi_n - \xi)$ is only relevant if the (lifted) sequence $\xi_n - \xi$ does not approach the tip of $C_{\mathfrak{p}}$ from inside $C_{\mathfrak{p}}$ but rather along its boundary. This picture of the topology of Ξ using cones is reminiscent to the oldest one, see [44], but we refer the reader to [3] for a proof in the present framework.

5.3. The Ellis semigroups of the systems $(\Xi, \mathbb{Z}^N)$ and $(\Omega_{\mathcal{L}}, \mathbb{R}^N)$

We now consider the Ellis semigroup $E(\Xi, \mathbb{Z}^N)$ of the dynamical system defined by the data $(\{A_i\}_{i\in I}, \Gamma, \Delta)$. We describe its topology, its monoid structure, and finally its action on Ξ.

Definition 5.8. A transformation type is a function $\mathfrak{t} \in \{+,-,0\}^I$ such that the cone

$$C'_{\mathfrak{t}} := \bigcap_{i\in I} H_i^{\mathfrak{t}(i)}$$

is non-empty. We denote by $\mathfrak{T}'$ the finite set of transformation types.

Note that the cones $C'_{\mathfrak{t}}$ are the constituents of a stratification of $\mathbb{R}^{N^\perp}$: If the $\mathfrak{t}(i)$ are all different from 0 then $C'_{\mathfrak{t}}$ is a connected component of $\mathbb{R}^{N^\perp} \setminus \bigcup_{i\in I} H_i^0$.

In general $C'_{\mathfrak{t}}$ is a connected component of $\bigcap_{i:\mathfrak{t}(i)=0} H_i^0 \setminus \bigcup_{i:\mathfrak{t}(i)\neq 0} H_i^0$. We denote by $\mathfrak{o} \in \{+,-,0\}^I$ the function which is constant equal to 0. Its cone is one point: $C'_{\mathfrak{o}} = \{0\}$. All cones are disjoint and so only $C'_{\mathfrak{o}}$ contains the origin.

Let $W_{\mathfrak{t}}$ be the connected component of $\overline{\mathbb{R}C'_{\mathfrak{t}} \cap \Gamma}$ containing 0. We call $C_{\mathfrak{t}} := W_{\mathfrak{t}} \cap C'_{\mathfrak{t}}$ the *effective* or *transformation* cone of $\mathfrak{t}$. It might be empty, as it is, for instance, if the intersection $\mathbb{R}C'_{\mathfrak{t}} \cap \Gamma$ is discrete but $\mathfrak{t} \neq \mathfrak{o}$. Let

$$\mathfrak{T} = \{\mathfrak{t} \in \mathfrak{T}' : C_{\mathfrak{t}} \neq \emptyset\}$$

be the set of effective transformation types. For $\mathfrak{t} \in \mathfrak{T}$ we consider $\mathbb{R}_{\mathfrak{t}}^{N^\perp} = \mathbb{R}C_{\mathfrak{t}} + \Gamma$ and its quotient $\mathbb{T}_{\mathfrak{t}}^{\perp} = \mathbb{R}_{\mathfrak{t}}^{N^\perp}/\Delta$. Note that $\Gamma/\Delta \subset \mathbb{T}_{\mathfrak{t}}^{\perp} \subset \mathbb{T}^{\perp}$.

Remark. The *complexity* of a Delone set $\mathcal{L}$ is the growth rate, as $R \to \infty$, of the number of translationally inequivalent sets of the form $B_R(x) \cap \mathcal{L}$: the complexity is α if this number grows like R^α. It is shown in [31] that, for an almost canonical cut-and-project set, the complexity α satisfies $N \leq \alpha \leq NN^\perp$. While maximal complexity ($\alpha = NN^\perp$) is generic, many of the familiar examples – the octagonal tilings, the Penrose tilings and their three-dimensional icosahedral generalisations, as well as the Danzer tilings – have minimal ($\alpha = N$) complexity. It is a feature of almost canonical projection sets of minimal complexity that $\Gamma \cap C'_{\mathfrak{t}}$ is dense in $C'_{\mathfrak{t}}$. For those systems the effective cone $C_{\mathfrak{t}}$ coincides with $C'_{\mathfrak{t}}$, $W_{\mathfrak{t}} = \bigcap_{i:\mathfrak{t}(i)=0} H_i^0$, and

$$\mathbb{T}_{\mathfrak{t}}^{\perp} = \left(\bigcap_{i:\mathfrak{t}(i)=0} H_i^0 + \Gamma \right) / \Delta.$$

Theorem 5.9 ([3]). *With the notation above, the operation*

$$(\mathfrak{t}\mathfrak{t}')(i) = \begin{cases} \mathfrak{t}(i) & \text{if } \mathfrak{t}(i) \neq 0 \\ \mathfrak{t}'(i) & \text{if } \mathfrak{t}(i) = 0 \end{cases}$$

defines a monoid structure[6] *on $\mathfrak{T}$ with $\mathfrak{t} = \mathfrak{o}$ as unit. The Ellis semigroup of $(\Xi, \mathbb{Z}^N)$ is isomorphic to the submonoid*

$$E(\Xi, \mathbb{Z}^N) \cong \bigcup_{\mathfrak{t} \in \mathfrak{T}} \mathbb{T}_{\mathfrak{t}}^{\perp} \times \{\mathfrak{t}\}$$

of $\mathbb{T}^\perp \times \mathfrak{T}$ equipped with the product

$$(\xi, \mathfrak{t})(\xi', \mathfrak{t}') = (\xi + \xi', \mathfrak{t}\mathfrak{t}').$$

Its action $E(\Xi, \mathbb{Z}^N) \times \Xi \to \Xi$ is given by

$$(\xi, \mathfrak{t}) \cdot (\xi', \mathfrak{p}) = (\xi + \xi', \mathfrak{p}') \text{ where } \mathfrak{p}'(i) = \begin{cases} \mathfrak{t}(i) & \text{if } i \in I(\xi + \xi') \text{ and } \mathfrak{t}(i) \neq 0 \\ \mathfrak{p}(i) & \text{if } i \in I(\xi + \xi') \text{ and } \mathfrak{t}(i) = 0 \\ \infty & \text{else} \end{cases}.$$

[6] We note a difference between this formula and the one in [2, 3] where the convention of [4] that the semigroup acts from the right is used.

The topology of $E(\Xi, \mathbb{Z}^N)$ is first countable and may thus be described in terms of convergence of sequences. A sequence $(\xi_n, \mathfrak{t}_n)_n$ in $\bigcup_{\mathfrak{t}\in\mathfrak{T}} \mathbb{T}_\mathfrak{t}^\perp \times \{\mathfrak{t}\}$ converges to $(\xi, \mathfrak{t})$ if and only if $\xi_n \to \xi$ in $\mathbb{T}^\perp$ and eventually $\xi_n - \xi \subset C_\mathfrak{t}$ and $T_{\overline{C_{\mathfrak{t}_n}}}(0) \subset T_{\overline{C_\mathfrak{t}}}(\xi_n - \xi)$.

Again the expressions $\xi_n - \xi \subset C_\mathfrak{t}$ and $T_{\overline{C_{\mathfrak{t}_n}}}(0) \subset T_{\overline{C_\mathfrak{t}}}(\xi_n - \xi)$ should be understood for large enough n so that we can lift $\xi_n - \xi$ into a small neighborhood of 0 in $\mathbb{R}^{N^\perp}$.

Remarks. The product on $\mathfrak{T}$ can be described geometrically with the help of the transformation cones. We do this below in the case of the octagonal tiling.

As it should be, the domain of $\mathfrak{p}'$ in the above formula is $I(\xi + \xi')$. Indeed, if $t(i) = 0$ then $\xi \in \mathbb{T}_\mathfrak{t} \subset H_i^0 + \Gamma$ and thus $i \in I(\xi + \xi')$ iff $i \in I(\xi')$. This implies that for $i \in I(\xi + \xi')$ with $t(i) = 0$ we must have $\mathfrak{p}(i) \neq \infty$.

Convergence of $(\xi_n, \mathfrak{t}_n)$ to $(\xi, \mathfrak{t})$ implies convergence of ξ_n to ξ in $\mathbb{T}^\perp$. Furthermore the copy of the acting group $\mathbb{Z}^N$ in $E(\Xi, \mathbb{Z}^N)$ is given by $\alpha_t = ([t], \mathfrak{o})$. As $T_{C_\mathfrak{o}}(0) = \{0\}$ we have that α_{t_n} converges to the transformation $(\xi, \mathfrak{t}) \in E(\Xi, \mathbb{Z}^N)$ if and only if $[t_n] \to \xi$ in $\mathbb{T}^\perp$ and eventually $[t_n] - \xi \in C_\mathfrak{t}$.

The Ellis semigroup of the continuous dynamical system $(\Omega_\mathcal{L}, \mathbb{R}^N)$ is just the suspension of $E(\Xi, \mathbb{Z}^N)$,

$$E(\Omega_\mathcal{L}, \mathbb{R}^N) \cong E(\Xi, \mathbb{Z}^N) \times_{\mathbb{Z}^N} \mathbb{R}^N.$$

The following theorem is thus the continuous version of Theorem 5.9.

Theorem 5.10 ([3]). *Consider the dynamical system $(\Omega_\mathcal{L}, \mathbb{R}^N)$ of an almost canonical cut-and-project set. There exists a finite monoid of idempotents $\mathfrak{T}$ which has a unique minimal left ideal $\mathfrak{T}_{min}$, and for each $\mathfrak{t} \in \mathfrak{T}$, a group $\mathbb{T}_\mathfrak{t}$ with $\mathbb{R}^N \subset \mathbb{T}_\mathfrak{t} \subset \mathbb{T} = \Omega_{\max}$ such that, algebraically,*

$$E(\Omega_\mathcal{L}, \mathbb{R}^N) \cong \bigcup_{\mathfrak{t}\in\mathfrak{T}} \mathbb{T}_\mathfrak{t} \times \{\mathfrak{t}\} \subset \mathbb{T} \times \mathfrak{T}$$

with semigroup law

$$(\xi, \mathfrak{t})(\xi, \mathfrak{t}') = (\xi + \xi', \mathfrak{t}\mathfrak{t}').$$

In particular the Ellis semigroup is a finite disjoint union of groups. For the unit $\mathfrak{o} \in \mathfrak{T}$ we have $\mathbb{T}_\mathfrak{o} = \mathbb{R}^N$ and for each minimal idempotent $\mathfrak{t} \in \mathfrak{T}$ we have $\mathbb{T}_\mathfrak{t} = \mathbb{T}$. Finally, the semigroup morphism induced by $\pi_{\max}$ is given by the projection onto the first factor.

When we say that $E(\Omega_\mathcal{L}, \mathbb{R}^N)$ is a finite disjoint union of groups we mean that the semigroup law restricted to the component $\mathbb{T}_\mathfrak{t} \times \{\mathfrak{t}\}$ is a group law, which follows here since $\mathfrak{t}\mathfrak{t} = \mathfrak{t}$ and so $([0], \mathfrak{t})$ is the neutral element in $\mathbb{T}_\mathfrak{t} \times \{\mathfrak{t}\}$. But this does not mean that $E(\Omega_\mathcal{L}, \mathbb{R}^N)$ is a group. $\mathfrak{T}$ is never a group. Moreover, as is the case for Lemma 5.4, the identification of $E(\Omega_\mathcal{L}, \mathbb{R}^N)$ as a submonoid of $\Omega_{\max} \times \mathfrak{T}$ does not respect the topology; in fact, the above theorem says nothing about the topology of $E(\Omega_\mathcal{L}, \mathbb{R}^N)$. The local nature of the topology, and the fact that it's first countable, can be got from Theorem 5.9.

In many cases the semigroup $\mathfrak{T}$ is very small, containing besides $\mathfrak{o}$ only minimal idempotents and so $E(\Omega_{\mathcal{L}}, \mathbb{R}^N) = \mathbb{T} \times \mathfrak{T}_{min} \cup \mathbb{R}^N$. These cases constitute the generic situation in [22] and correspond to cut-and-project sets with *maximal* complexity [31]. On the opposite side, the almost canonical cut-and-project sets with minimal complexity have the largest possible $\mathfrak{T}$ (see the remark preceding Theorem 5.9). Less complexity seems to make the Ellis semigroup richer! Almost canonical cut-and-project sets with minimal complexity share also the property that their rational cohomology groups are finitely generated [31, 22]. All Delone sets coming from primitive substitutions and, more generally, all linearly repetitive tilings, have minimal complexity. This can be rather directly inferred from [49] and is discussed in some detail in [7]. There it is also shown that pure point spectrum implies zero entropy (i.e., sub-exponential complexity) for general uniquely ergodic systems.

5.4. Example of the octagonal tiling

The octagonal tiling has dimension and codimension 2. Its window W is a regular octahedron which is the projection of the unit cube for $\tilde{L} = \mathbb{Z}^4$ onto $\mathbb{R}^{N^\perp} = \mathbb{R}^2$. Γ is the lattice generated by the difference vectors between corners of the octahedron. Of the eight affine hyperplanes spanned by the sides of the octahedron, only four are independent modulo Γ so we only need four lines H_i^0 and can take $a_i = 0$ to describe the affine hyperplanes furnishing the input data of the dynamical system. These four lines form a regular 8-star in $\mathbb{R}^2$. We number them so that H_1^0 and H_3^0 are orthogonal and hence also H_2^0 and H_4^0.

It turns out that we have 8 possible cut types [22]: If $\xi \in \Gamma/\Delta$ then $I(\xi) = \{1,2,3,4\}$, i.e., ξ lies on four different affine hyperplanes. If $\xi \in (\frac{e_1+e_3}{2} + \Gamma)/\Delta$ then $I(\xi) = \{2,4\}$ and, if $\xi \in (\frac{e_2+e_4}{2} + \Gamma)/\Delta$ then $I(\xi) = \{1,3\}$. If $\xi \in (H_i + \Gamma)/\Delta$ but it lies not in the above sets then the cut type is $I(\xi) = \{i\}$. Here $i = 1,2,3,4$ so these yield four cut types. Finally, the cut type of all other points is $I(\xi) = \emptyset$.

It follows that we have 25 different point types: Cut type $\{1,2,3,4\}$ allows 8 different point types whose cones correspond to the 8 cones with opening angle of 45 degree and boundary contained in $H_1^0 \cup H_2^0 \cup H_3^0 \cup H_4^0$. Cut type $\{1,3\}$ and $\{2,4\}$ allow each for four point types which correspond to the 4 cones with operning angle of 90 degree and boundary contained in $H_1^0 \cup H_3^0$ and $H_2^0 \cup H_4^0$, respectively. Cut type $\{i\}$ allows for two point types corresponding to the two open half-spaces bounded by H_i^0. Finally, if $I(\xi) = \emptyset$ we have a single point type $I = \underline{\infty}$ with cone all of $\mathbb{R}^2$.

An element $(\xi, \mathfrak{p}) \in \Xi$ corresponds to a tiling. The elements with point type $\underline{\infty}$ correspond to the non-singular tilings. The other elements correspond to tilings which have worms. A worm[7] is a configuration of tiles along a line, in the octagonal tiling made of squares and rhombi, which may occur in two different orientations. More precisely, three rhombi fill a hexagon, and they can do this in two different ways. In the worm one or the other way is realised coherently for all hexagons and changing the way is referred to as flipping the worm. Now whenever $i \in I(\xi)$ then

[7] In the Penrose tiling these worms are referred to as Conway worms [27].

the tiling corresponding to $(\xi, \mathfrak{p}) \in \Xi$ contains a worm in a direction determined by i. So depending on the cut type the tiling has one, two, or four worms in one, two, or four directions, respectively, and the point type corresponds precisely to the information in which way the worms are flipped. The first coordinate ξ carries the information on where the worms cross. This describes the space Ξ. The $\mathbb{Z}^2$-action (N=2) is given by translation of the first variable which amounts to translation of the tiling.

We now describe $E(\Xi, \mathbb{Z}^2)$. The octagonal tiling has the property that $C'_{\mathfrak{t}} \cap \Lambda$ is dense in $C'_{\mathfrak{t}}$ provided $\mathfrak{t} \neq \mathfrak{o}$ which implies that all cones coincide with their effective cones. In the octagonal case these cones have dimension 2, or 1, or 0, the latter only for $C_{\mathfrak{o}}$. The two-dimensional cones coincide with the 8 cones of the point types which have an opening angle of 45 degree. They are associated to the minimal idempotents and their corresponding group $\mathbb{T}^{\perp}_{\mathfrak{t}}$ is equal to $\mathbb{T}^{\perp}$. Furthermore there are eight one-dimensional cones corresponding to half-lines, more precisely for each i one or the other half of H^0_i. Their corresponding group is $(H^0_i + \Gamma)/\Delta$. Finally there is the zero-dimensional cone $C_{\mathfrak{o}} = \{0\}$ whose group is $\Gamma/\Delta = \mathbb{Z}^2$.

The product on $\mathfrak{T}$ can now be described geometrically by means of the associated cones. If $C^{(2)}_1$ is a two-dimensional cone and $C^{(*)}_2$ any other cone we have the left domination rule $C^{(2)}_1 C^{(*)}_2 = C^{(2)}_1$. If $C^{(1)}_1$ is a one-dimensional cone and $C^{(2)}_2$ a two-dimensional one then the result is the two-dimensional cone $C^{(1)}_1 C^{(2)}_2 = C^{(2)}_3$ which can be described as "bringing $C^{(2)}_2$ alongside": $C^{(1)}_1$ is a half-line and $C^{(2)}_2$ an open cone which is on a distinguished side of the line to which the half-line belongs. $C^{(2)}_3$ is then the open cone which is on the same side as $C^{(2)}_2$ and moreover contains $C^{(1)}_1$ in its boundary. Almost the same sort of rule applies to the product of two one-dimensional cones $C^{(1)}_1$ and $C^{(1)}_2$. If these correspond to half-lines which are not parallel then $C^{(1)}_1 C^{(1)}_2 = C^{(2)}_3$ where $C^{(2)}_3$ is the open cone which is on the same side of $C^{(1)}_1$ as $C^{(2)}_1$ and contains $C^{(1)}_1$ in its boundary. If however $C^{(1)}_1$ and $C^{(1)}_2$ are parallel then the left domination rule applies: $C^{(1)}_1 C^{(1)}_2 = C^{(1)}_1$. Finally, there is only one cone of dimension 0, this cone corresponds to the unit.

As we have already mentionned, the transformation cones are all disjoint. But the inclusion of a cone in the (Euclidean) closure of another cone has an algebraic interpretation. Indeed, if $C_{\mathfrak{t}} \subset \overline{C_{\mathfrak{t}'}}$ then $\mathfrak{t}$ and $\mathfrak{t}'$ satisfy $\mathfrak{t}\mathfrak{t}' = \mathfrak{t}'\mathfrak{t} = \mathfrak{t}'$ which corresponds exactly to the definition of the natural order on the set of idempotents of a semigroup: $\mathfrak{t} \geq \mathfrak{t}'$ whenever $\mathfrak{t}\mathfrak{t}' = \mathfrak{t}'\mathfrak{t} = \mathfrak{t}'$.

The action of E on Ξ is more complicated, as the second coordinate depends on the point type. But it simplifies in the particular case of an idempotent as follows: $([0], \mathfrak{t}) \cdot (\xi, \mathfrak{p}) = (\xi, \mathfrak{p}')$ with $\operatorname{dom} \mathfrak{p}' = \operatorname{dom} \mathfrak{p}$ and

$$\mathfrak{p}'(i) = \begin{cases} \mathfrak{t}(i) & \text{if } \mathfrak{t}(i) \neq 0 \\ \mathfrak{p}(i) & \text{if } \mathfrak{t}(i) = 0 \end{cases} .$$

This can again be geometrically described in terms of the associated cones. The action of $\mathfrak{t}$ is like "bringing the point cone along": the point cone $C_{\mathfrak{p}'}$ is the cone

which has the same point type as $\mathfrak{p}$, contains the cone $C_{\mathfrak{t}}$ in its closure, and, provided this does not specify it uniquely, lies on the same side as $C_{\mathfrak{p}}$ w.r.t. the vector space spanned by $C_{\mathfrak{t}}$. In particular, point type $\underline{\infty}$ is invariant under the action of a transformation type. The effect of the action of the idempotent $([0], \mathfrak{t})$ of the Ellis semigroup on the tiling corresponding to $(\xi, \mathfrak{p})$ is thus as follows: If the tiling is non-singular then it acts like the identity. If the tiling is singular then the action is to flip the worms into (or to leave them in) the position which is encoded by $\mathfrak{p}'$.

If the acting element $(\xi, \mathfrak{t})$ is not an idempotent, then things may become more complicated, except if $\xi \in \Gamma$ in which case the point type of $\xi + \xi'$ agrees again with that of ξ' and we have the combination of a translation of the tiling with a flip of worms. Otherwise the formula for the action given in Theorem 5.9 takes into account that the domain of $\mathfrak{p}'$ coincides with the point type of $\xi + \xi'$ and tilings may be mapped to tilings with distinct worm configurations.

References

[1] J.E. Anderson and I.F. Putnam, *Topological invariants for substitution tilings and their associated C^*-algebras*, Ergodic Th. & Dynam. Syst. **18** (1998), 509–537.

[2] J.-B. Aujogue, *On embedding of repetitive Meyer multiple sets into model multiple sets*, to appear in Ergodic Theory and Dynamical Systems, arXiv:1405.3013.

[3] J.-B. Aujogue, *Ellis enveloping semigroup for almost canonical model sets of an Euclidean space*, to appear in Algebraic and Geometric Topology, arXiv:1305.0879v2

[4] J. Auslander, *Minimal Flows and their Extensions*, North-Holland Mathematical Studies 153, Elsevier 1988.

[5] M. Baake and D. Lenz, *Dynamical systems on translation bounded measures: Pure point dynamical and diffraction spectra*, Ergodic Th. & Dynam. Syst. **24** (2004), 1867–1893; math.DS/0302061.

[6] M. Baake, D. Lenz and R.V. Moody, *A characterization of model sets by dynamical systems*, Ergodic Theory Dynam. Systems **27** (2007), 341–382.

[7] M. Baake, D. Lenz and C. Richard, *Pure point diffraction implies zero entropy for Delone sets with uniform cluster frequencies*, Lett. Math. Phys. **82**, (2007), 61–77.

[8] M. Baake, R.V. Moody, P. Pleasants, *Diffraction from visible lattice points and k-th power free integers* Discr. Math. **221** (1999), 3–42.

[9] M. Baake and R.V. Moody (eds), *Directions in Mathematical Quasicrystals*, CRM Monograph Series, vol. 13, AMS, Providence, RI (2000).

[10] M. Baake and R.V. Moody, *Weighted Dirac combs with pure point diffraction*, J. Reine Angew. Math. (Crelle) **573** (2004), 61–94; math.MG/0203030.

[11] M. Barge, J. Kellendonk, *Proximality and pure point spectrum for tiling dynamical systems*, Michigan Math. J. **64**, no. 4, (2013), 793–822.

[12] M. Barge, C. Olimb, *Asymptotic structure in substitution tiling spaces*, Ergod. Th. and Dynam. Sys.**34**, No.1, (2014), 55–94.

[13] J. Bellissard, D. Herrmann, M. Zarrouati, *Hull of Aperiodic Solids and Gap Labelling Theorems*, in [9], 207–259.

[14] Kenneth S. Brown "Semigroups, rings, and Markov chains." Journal of Theoretical Probability 13.3 (2000): 871–938.

[15] J.P. Clay, Proximity relations in transformation groups, *Trans. Amer. Math. Soc.* **108** (1963), 88–96.

[16] M. Denker, C. Grillenberger and K. Sigmund, *Ergodic Theory on Compact Spaces*, Lecture Notes in Mathematics, vol. 527, Springer, Berlin (1976).

[17] S. Dworkin, *Spectral theory and X-ray diffraction*, J. Math. Phys. 34(1993), 2965–2967.

[18] M. Einsiedler and T. Ward, *Ergodic theory with a view towards number theory*, vol. 259 of *Graduate Texts in Mathematics*. Springer-Verlag London, Ltd., London, 2011.

[19] R. Ellis, *Lectures on topological dynamics*, W.A. Benjamin, Inc., New York, 1969.

[20] R. Ellis, *The Veech structure theorem*, Transactions of the American Mathematical Society 186 (1973), 203–218.

[21] R. Ellis and M. Nerurkar, *Weakly almost periodic flows*, Trans. Amer. Math. Soc. 313 (1989), 103–119.

[22] Forrest, A.H.; Hunton, J.R.; Kellendonk, J.: Topological invariants for projection method patterns. Mem. Amer. Math. Soc. **758** (2002).

[23] H. Furstenberg, *The structure of distal flows*, Amer. J. of Math., **85** (1963), 477–515.

[24] E. Glasner, *On tame dynamical systems*, Colloq. Math., 105 (2006), 283–295.

[25] E. Glasner and M. Megrelishvili, *Hereditarily non-sensitive dynamical systems and linear representations*, Colloq. Math., 104 (2006), no. 2, 223–283.

[26] J.-B. Gouéré, *Quasicrystals and almost periodicity*, Commun. Math. Phys. **255** (2005), 655–681.

[27] B. Grünbaum, G.C. Shephard. *Tilings and patterns*. Freeman, 1987.

[28] A. Hof, *On diffraction by aperiodic structures*, Commun. Math. Phys. **169** (1995), 25–43.

[29] A. Hof, *Diffraction by aperiodic structures*, in: [9], pp. 239–268.

[30] B. Host, *Valeurs propres des systèmes dynamiques définis par des substitutions de longueur variable*, Ergodic Th. Dynam. Sys. **6** (1986), 529–540.

[31] A. Julien. Complexity and Cohomology for cut-and-projection tilings. *Ergodic Theory and Dynamical Systems* 30 (2010) 489–523.

[32] J. Kellendonk, D. Lenz, *Equicontinuous Delone dynamical systems*, Canadian Journal of Mathematics **65** (2013), 149–170.

[33] J. Kellendonk, L. Sadun, *Meyer sets, topological eigenvalues, and Cantor fiber bundles*, Journal of the London Mathematical Society (2013): jdt062.

[34] J. Kelly, General Topology, Springer-Verlag, Berlin-New York (1991).

[35] R. Kenyon, *Self-similar Tilings*, Ph.D. Thesis, Princeton University, 1990.

[36] R. Kenyon and B. Solomyak, *On the characterization of expansion maps for self-affine tilings*, Discrete Comp. Geom. **43** (2010), 577–593.

[37] P. Kurka, Topological and symbolic dynamics. Société Mathématique de France, (2003).

[38] J. Kwapisz, *Expansions of self-affine tilings are integral Perron* (*the non-diagonalizable case*), Preprint, 2013.

[39] J.C. Lagarias, *Meyer's concept of quasicrystal and quasiregular sets*, Commun. Math. Phys. **179** (1996), 365–376.

[40] J.C. Lagarias, *Geometric models for quasicrystals I. Delone sets of finite type*, Discrete Comput. Geom. **21** (1999), 161–191.

[41] J. Lagarias, *Mathematical quasicrystals and the problem of diffraction*, in: [9], 61–93.

[42] J. Lagarias, P.A.B. Pleasants, *Repetitive Delone sets and quasicrystals*, Ergodic Theory Dynam. Systems **23** (2003), 831–867.

[43] J. Lagarias, P.A.B. Pleasants, *Local complexity of Delone sets and crystallinity*, Canad. Math. Bull. **45** (2002), 634–652

[44] T.T.Q. Le, *Local Rules for Quasiperiodic Tilings*, In: The Mathematics of Long Range Aperiodic Order, 331–366. Kluwer 1997.

[45] J.-Y. Lee, *Substitution Delone sets with pure point spectrum are inter-model sets*, J. Geom. Phys. **57** (2007), 2263–2285.

[46] J.-Y. Lee, R.V. Moody, B. Solomyak, *Pure point dynamical and diffraction spectra.* Annales Henri Poincaré. Vol. 3. No. 5, 2002.

[47] J.-Y. Lee and B. Solomyak, *Pisot family self-affine tilings, discrete spectrum, and the Meyer property*, Discrete Contin. Dyn. Syst., **32**, no. 3, (2012), 935–959.

[48] J.-Y. Lee and B. Solomyak, *Pure point diffractive substitution Delone sets have the Meyer property*, Discrete Comp. Geom. **34** (2008), 319–338.

[49] D. Lenz, *Aperiodic linearly repetitive Delone sets are densely repetitive*, Discrete Comput. Geom. **31**, (2004), 323–326.

[50] D. Lenz, *Continuity of eigenfunctions of uniquely ergodic dynamical systems and intensity of Bragg peaks*, Commun. Math. Phys. **287** (2009), 225–258.

[51] D. Lenz, C. Richard, *Pure point diffraction and cut and project schemes for measures: The smooth case*, Math. Z. **256** (2007), 347–378.

[52] D. Lenz and P. Stollmann, *Delone dynamical systems and associated random operators*, OAMP Proceedings, Constanta 2001; J.-M. Combes, J. Cuntz, G.A. Elliott, G. Nenciu, H. Siedentop, S. Stratila (eds.), Theta Foundation, ISBN 973-85432-2-3.

[53] D. Lenz, N. Strungaru, *Pure point spectrum for measure dyamical systems on locally compact Abelian groups*, J. Math. Pures Appl. **92** (2009), 323–341.

[54] Y. Meyer, *Algebraic Numbers and Harmonic Analysis*, North-Holland, Amsterdam (1972).

[55] R.V. Moody, ed., *The Mathematics of Long-Range Aperiodic Order*, NATO ASI Series C 489, Kluwer, Dordrecht (1997).

[56] R.V. Moody, *Meyer sets and their duals*, in: [55], 403–441.

[57] R.V. Moody, *Model sets: A Survey*, in: *From Quasicrystals to More Complex Systems*, eds. F. Axel, F. Dénoyer and J.P. Gazeau, EDP Sciences, Les Ulis, and Springer, Berlin (2000), pp. 145–166.

[58] R.V. Moody, *Mathematical quasicrystals: a tale of two topologies.* XIVth International Congress of Mathematical Physics, ed. J.-C. Zambrini, World Scientific, Singapore. 2005.

[59] J. Patera (ed.), *Quasicrystals and Discrete Geometry*, Fields Institute Monographs, vol. 10, AMS, Providence, RI (1998).

[60] M. Petrich and N.R. Reilly. *Completely regular semigroups*. New York: Wiley, 1999.

[61] M. Quéffélec, Substitution dynamical systems – spectral analysis. Lecture Notes in Mathematics, 1294. Springer-Verlag, Berlin, 1987.

[62] E.A. Robinson, *On uniform convergence in the Wiener-Wintner theorem*, J. London Math. Soc. **49** (1994), 493–501.

[63] E.A. Robinson, *The dynamical properties of Penrose tilings*, Trans. Amer. Math. Soc. **348** (1996), 4447–4464.

[64] E.A. Robinson, *A Halmos von Neumann Theorem for model sets, and almost automorphic dynamical systems* in: Recent progress in dynamics, MSRI publications **57**, (2007), 243–272.

[65] E.A. Robinson, Jr., "Symbolic dynamics and tilings of $\mathbb{R}^d$," in *Symbolic dynamics and its applications*, vol. 60 of *Proc. Sympos. Appl. Math.*, pp. 81–119, Providence, RI: Amer. Math. Soc., 2004.

[66] D.J. Rudolph, Rectangular tilings of R^n and free R^n-actions, *Dynamical systems* (College Park, MD, 198–87), Springer, Berlin, 1988, pp. 653–688.

[67] L. Sadun and R.F. Williams, *Tiling spaces are Cantor set fiber bundles*, Ergodic Theory and Dynamical Systems **23** (2003), 307–316.

[68] M. Schlottmann, Generalized model sets and dynamical systems, in:[9], pp. 143–159.

[69] B. Sing, *Pisot Substitutions and Beyond*, PhD thesis, Universität Bielefeld (2007).

[70] B. Solomyak, *Dynamics of self-similar tilings*, Ergodic Th. & Dynam. Syst. **17** (1997) 695–738; Erratum: Ergodic Th. & Dynam. Syst. **19** (1999) 1685.

[71] B. Solomyak, *Spectrum of dynamical systems arising from Delone sets*, in: [59] pp. 265–275.

[72] B. Solomyak, *Nonperiodicity implies unique composition for self-similar translationally finite tilings*, Discrete Comput. Geometry **20** (1998), 265–279.

[73] B. Solomyak, *Eigenfunctions for substitution tiling systems*, Advanced Studies in Pure Mathematics **49** (2007), 433–454, (Proceedings of the International Conference on Probability and Number Theory in Kanazawa, Japan, June 2005).

[74] N. Strungaru, *Almost periodic measures and long-range order in Meyer sets.* Discrete Comput. Geom. **33** (2005), 483–505.

[75] W.A. Veech, *Point-distal flows*, American Journal of Mathematics **92**, no. 1, (1970), 205–242.

[76] P. Walters, *An Introduction to Ergodic Theory*, Springer, New York (1982).

Jean-Baptiste Aujogue
Universidad de Santiago de Chile
Dep. de Matemáticas, Fac. de Ciencia
Aladema 3363, Estación Central
Santiago, Chile

Marcy Barge
Department of Mathematical Sciences
Montana State University
P.O. Box 172400
Bozeman, MT 59717-2400, USA

Johannes Kellendonk
Université Claude Bernard Lyon 1
CNRS, UMR 5208 Institut Camille Jordan
F-69622 Villeurbanne cedex, France

Daniel Lenz
Friedrich-Schiller-Universität Jena
Fakultät für Mathematik und Informatik
Mathematisches Institut
Ernst-Abbe-Platz 2
D-07743 Jena, BRD

Progress in Mathematics, Vol. 309, 195–222
© 2015 Springer Basel

Linearly Repetitive Delone Sets

José Aliste-Prieto, Daniel Coronel, María Isabel Cortez,
Fabien Durand and Samuel Petite

Abstract. Linearly repetitive Delone sets are the simplest aperiodic repetitive Delone sets of the Euclidean space, e.g. any self similar Delone set is linearly repetitive. We present here some combinatorial, ergodic and mixing properties of their associated dynamical systems. We also give a characterization of such sets via the patch frequencies. Finally, we explain why a linearly repetitive Delone set is the image of a lattice by a bi-Lipschitz map.

Mathematics Subject Classification (2010). 37B50.

Keywords. Delone sets, tiling systems, linearly repetitive, Voronoï cell.

1. History and motivations

The notion of *linearly recurrent subshift* has been introduced in [Du, DHS] to study the relations between substitutive dynamical systems and stationary dimension groups. In an independent way, the similar notion of *linearly repetitive Delone sets* of the Euclidean space $\mathbb{R}^d$ appears in [LP1]. For a Delone set X of $\mathbb{R}^d$, the *repetitivity function* $M_X(R)$ is the least M (possibly infinite) such that every closed ball B of radius M intersected with X contains a translated copy of any patch with diameter smaller than $2R$.

A Delone set X is said *linearly repetitive* if there exists a constant L such that $M_X(R) < LR$ for all $R > 0$. Observe that we can assume that the constant L is greater than 1. According to the following theorem, the slowest growth for the repetitivity function of an aperiodic Delone set is linear.

Theorem 1 ([LP1, Thm. 2.3]). *Let $d \geq 1$. There exists a constant $c(d) > 0$ such that for any Delone set X of $\mathbb{R}^d$ such that*

$$M_X(R) < c(d)R \qquad \text{for some } R > 0,$$

then X has a non-zero period.

Financial support from the ANR SUBTILE 0879. This work is part of the program MathAmSud DYSTIL 12Math-02.

Even more, if for some R, $M_X(R) < \frac{4}{3}R$, then the Delone set X is a *crystal*, *i.e.*, has d independent periods ([LP1, Thm. 2.2]).

The classical examples of aperiodic Delone systems, e.g., the ones arising from substitution, are linearly repetitive.

Lemma 2 ([So2, Lem. 2.3]). *A primitive self similar tiling is linearly repetitive.*

In many senses that we will not specify, the family of linearly repetitive Delone sets is small inside the family of all the Delone sets of the Euclidean space $\mathbb{R}^d$. For instance, in the class of Sturmian subshifts, several authors [MH, Du1, Du, LP2] show the following result.

Proposition 3. *The Sturmian subshift associated to an irrational number α is linearly recurrent if and only if the coefficients of the continued fraction of α are bounded.*

Let us recall that for the standard topology, the set of numbers with bounded continued fraction are badly approximable by rational numbers. It is known that they form a Baire meager set, with 0 Lebesgue measure but with Hausdorff dimension 1.

As we shall see, the linearly repetitive Delone sets possess many rigid properties. In the next section we present some combinatorial properties of these sets. For instance, their complexity appears to be the slowest possible among all the aperiodic repetitive Delone sets. Section 3 is devoted to the structure of the hull of an aperiodic linearly repetitive Delone set. A tower system with uniform bound is described. We deduce from this the main properties of the system. We focus in Section 4 on the ergodic properties of dynamical systems associated to linearly repetitive Delone sets. They are strictly ergodic (i.e., each patch appears with a frequency). But they are not wild since they are never measurably mixing. They satisfy also a subbaditive ergodic theorem. We present a characterization of the linear repetitivity by using a bound on the frequencies of the occurrences of the patches. The dynamical factors of these systems are studied in Section 5. They admit as factors just a finite number of non conjugate aperiodic Delone systems. We give also a characterization of their continuous and measurable eigenvalues by studying cohomological equations. The last section concerns the deformation of linearly repetitive Delone sets: each one is the image through a Lipschitz map of a lattice in $\mathbb{R}^d$.

2. Combinatorial properties

In this section we give the basic definitions and combinatorial properties concerning linearly repetitive Delone sets of $\mathbb{R}^d$. Most of these properties are obvious for self-similar tilings. Recall that a set $X \subset \mathbb{R}^d$, with $d \geq 1$, is a (r_X, R_X)-Delone set (or a Delone set for short) if it is a discrete subset of the Euclidean space $\mathbb{R}^d$, with the following properties:

1. *Uniform discreteness*: each open ball of radius $r_X > 0$ in $\mathbb{R}^d$ contains at most one point of X.
2. *Relative density*: each closed ball of radius R_X in $\mathbb{R}^d$ contains at least one point of X.

A classical example is given by the lattice $\mathbb{Z}^d$ of points with integer coefficients is a Delone set. But notice also that the image of any Delone set by a bi-Lipschitz map of $\mathbb{R}^d$ provides a Delone set. We denote by $B_R(x)$ the Euclidean closed ball of radius $R > 0$ centered at the point $x \in \mathbb{R}^d$.

2.1. Return vectors to a patch

Let X be a (r_X, R_X)-Delone set. A *R-patch* is a set of the kind $\mathtt{P} = X \cap B_R(x)$ centered at some point $x \in X$ and for some $R > R_X$[1]. In the rest of this paper we assume that all the Delone sets have *finite local complexity*, that is for any real $R > 0$ there is a finite number of R-patches, up to translations. This is actually equivalent to the fact $X - X$ is a discrete subset of $\mathbb{R}^d$ [La].

For a R-patch $\mathtt{P}$, we define the set

$$\mathcal{R}_{\mathtt{P}}(X) = \{v \in \mathbb{R}^d : \mathtt{P} + v \text{ is a } R\text{-patch of } X\}.$$

It is called the set of *return vectors* to $\mathtt{P}$. For a fixed center $x_{\mathtt{P}}$ of $\mathtt{P}$, any point in $\mathcal{R}_{\mathtt{P}}(X) + x_{\mathtt{P}} =: X_{\mathtt{P}}$ is an *occurrence* of the patch $\mathtt{P}$.

Observe that the null vector 0 always belongs to $\mathcal{R}_{\mathtt{P}}(X)$. It is straightforward to check that $X_{\mathtt{P}}$ is a Delone set when X is linearly repetitive (see definition in the introduction). Furthermore, $X_{\mathtt{P}}$ has finite local complexity because $X_{\mathtt{P}} - X_{\mathtt{P}} \subset X - X$.

When X is aperiodic and linearly repetitive with constant L, there are uniform bounds on the constants $r_{X_{\mathtt{P}}}$ and $R_{X_{\mathtt{P}}}$ associated to the Delone set $X_{\mathtt{P}}$. The following lemma shows that two occurrences of a patch can not be too close. The proof can be found in [Le, Lem. 2.1] and in [So2, Du1].

Lemma 4. *Let X be a linearly repetitive aperiodic Delone set with constant $L > 1$. Then, for every patch $\mathtt{P} = X \cap B_R(x)$ with $x \in X$, $R > 0$, we have*

$$\frac{R}{L+1} \leq r_{X_P} \leq R_{X_P} \leq LR.$$

Proof. By contradiction: let us assume there exist $x \neq y \in X$ with

$$(X \cap B_R(x)) - x = X \cap B_R(y) - y$$

and

$$r_X \leq \|x - y\| < \frac{R}{(L+1)}.$$

Then for any point z' in $B_R(x) \cap X$, we have $z' + (y - x) \in X$. For any $z \in X$, the set $X \cap B_R(x)$ contains a translated copy centered in $z' \in X \cap B_R(x)$ of the patch $B_{\frac{R}{L+1}}(z) \cap X$. Thus $z' + (y - x) \in X \cap B_{\frac{R}{L+1}}(z')$ and finally $z + (y - x) \in X$ and

[1]Note: a given patch may be defined by several centers x and radius R. So when we consider a R-patch $\mathtt{P}$, we choose a center $x_{\mathtt{P}}$ and a radius R.

so $X + (y - x) \subset X$. In a similar way we obtain $X + (x - y) \subset X$, so that finally we get $X + x - y = X$ contradicting the aperiodicity of X. □

This repulsion property on the occurrences of patches has several consequences on the combinatorics of the Delone set X.

First of all on the complexity. Let us denote $N_X(R)$ the number of different R-patches $B_R(x) \cap X$ with $x \in X$, up to translation. Since any ball of radius $M_X(R)$ contains the centers of occurrences of any R-patch, we easily deduce that $N_X(R)^{\frac{1}{d}} = O(M_X(R))$ as $R \to \infty$ (see [LP2]).

Lemma 5 ([Le, Lem. 2.2]). *Let X be an aperiodic linearly repetitive Delone set. Then*

$$\liminf_{R \to +\infty} \frac{N_X(R)}{R^d} > 0.$$

From this, we conclude that for an aperiodic linearly repetitive Delone set $M_X(R) = O(N_X(R)^{\frac{1}{d}})$ as $R \to \infty$.

Proof. As X is relatively dense, there exist constants $\lambda_1 > 0$ and $R_1 > 0$ such that

$$\sharp(X \cap B_R(x)) \geq \lambda_1 R^d \qquad \text{for any } x \in X,\ R \geq R_1.$$

By the previous lemma all the patches $(X - x) \cap B_R(0)$ for $x \in X \cap B_{\frac{R}{3(L+1)}}(0)$ are pairwise different. Thus for any $R \geq 3(L+1)R_1$, we have

$$N_X(R) \geq \sharp(X \cap B_{\frac{R}{3(L+1)}}(0)) \geq \lambda_1 \left(\frac{R}{3(L+1)}\right)^d,$$

that gives us the result. □

Another property is on the hierarchical structure of the linearly repetitive Delone sets, that is quite simple: for any size $R > 0$, it is possible to decompose the Delone set into big patches (each one containing a R-patch), so that the number of these patches, up to translations, is independent of the size R. To be more precise, we need the notion of *Voronoï cell of a patch*. For a (r_X, R_X)-Delone set X, the *Voronoï cell* V_x of a point $x \in X$ is the set

$$V_x = \big\{ y \in \mathbb{R}^d : \|y - x\| \leq \|y - x'\|, \forall x' \in X \big\}.$$

It is then direct to check that any Voronoï cell V_x is a convex polyhedra, its diameter is smaller or equal to $2R_X$ and it contains the ball $B_{\frac{r_X}{2}}(x)$. Moreover when the Delone set X is of finite local complexity, the collection of Voronoï cells $\{V_x\}_{x \in X}$ forms a tiling of $\mathbb{R}^d$ of finite local complexity.

For a patch R-patch $\mathtt{P} = B_R(x_\mathtt{P}) \cap X$ of a repetitive Delone set X, we denote by $V_{\mathtt{P},x}$ the Voronoï cell associated to the Delone set $X_\mathtt{P}$ and an occurrence $x \in X_\mathtt{P}$. Notice that the Voronoï cell associated to the set of return vectors $\mathcal{R}_\mathtt{P}(X)$ and a return vector $v \in \mathcal{R}_\mathtt{P}(X)$, is the Voronoï cell of the occurrence $x_\mathtt{P} + v \in X_\mathtt{P}$ translated by the vector $-x_\mathtt{P}$.

It follows by Lemma 4 that for an aperiodic linearly repetitive Delone set with constant L, for any R-patch $\mathtt{P}$,

$$\operatorname{diam} V_{\mathtt{P},x} \le 2LR, \qquad B_{\frac{R}{2(L+1)}}(x) \subset V_{\mathtt{P},x}, \qquad \text{for any } x \in X_{\mathtt{P}}. \tag{2.1}$$

Lemma 6 ([CDP, Lem. 11]). *Let X be an aperiodic linearly repetitive Delone set with constant L. There exists an explicit positive constant $c(L)$ such that for every $R > 0$ and every R-patch $\mathtt{P} = X \cap B_R(x)$, the collection $\{X \cap V_{P,x} : x \in X_P\}$ contains at most $c(L)$ elements up to translation.*

Observe here that the bound, explicit in the proof, does not depend on the combinatorics of X but just on the constant of repetitivity.

Proof. Let us consider B the union of Voronoï cells $V_{\mathtt{P},x}$, $x \in X_{\mathtt{P}}$ that intersects the ball $B_{L^2R}(0)$. We have then

$$B_{L^2R}(0) \subset B \subset B_{L^2R+2LR}(0).$$

By linear repetitivity, $B \cap X$ contains a translated copy of any patch of the kind $X \cap V_{\mathtt{P},x}$ with $x \in X_{\mathtt{P}}$. Since any Voronoï cell contains a ball of radius $\frac{R}{2(L+1)}$, the number of patches in $B \cap X$ of the kind $X \cap V_{\mathtt{P},x}$ with $x \in X_{\mathtt{P}}$ is smaller than

$$\frac{\operatorname{vol} B_{RL(L+2)}(0)}{\operatorname{vol} B_{\frac{R}{2(L+1)}}(0)} \le (2L(L+2)^2)^d = c(L). \qquad \square$$

Even stronger, the next lemma gives for an aperiodic linearly repetitive Delone set, a uniform bound (in R) on the number of occurrences of a patch inside a ball of radius KR.

Lemma 7. *Let X be an aperiodic linearly repetitive Delone set with constant $L \ge 1$, and let $K \ge L$. Then for any R-patch $\mathtt{P}$ of X and any point $y \in \mathbb{R}^d$,*

$$\sharp\{v \in \mathbb{R}^d; \mathtt{P} - v \subset B_{KR}(y) \cap X\} \le 12^d K^d L^d.$$

Proof. Let B be the union of all the Voronoï cells $V_{\mathtt{P},x}$, $x \in X_{\mathtt{P}}$ that intersect the ball $B_{KR}(y)$. It follows that

$$B \subset B_{KR+2LR}(y).$$

By Lemma 4, the sets $B_{\frac{R}{2(L+1)}}(z)$, where the points $z \in B_{KR}(y) \cap X_{\mathtt{P}}$ are occurrences of $\mathtt{P}$, are pairwise disjoint and are included in B. Then it follows that

$$\sharp\{v \in \mathbb{R}^d; \mathtt{P} - v \subset B\} \le \frac{\operatorname{vol}(B)}{\operatorname{vol} B_{\frac{R}{2(L+1)}}(0)} \le 2^d(K+2L)^d(L+1)^d,$$

that gives us the result. $\square$

Here again, observe that the bound depends just on the repetitivity constant.

3. Structure of the hull of a linearly repetitive Delone set

3.1. Background on solenoids, boxes

In this section, we will see the specific geometrical structure of the associated hull Ω of an aperiodic repetitive Delone set. We recall here, from [BBG, BG], the local structure of this space.

3.1.1. Local transversals and return vectors. Let $(\Omega, \mathbb{R}^d)$ be an aperiodic minimal Delone system. The *canonical transversal* of Ω is the set Ω^0 composed of all Delone sets in Ω that contain the origin 0. This terminology is motivated by the fact that if Y is in Ω^0, then every small translation of Y will not be in Ω^0. A *cylinder* in Ω is a set of the form

$$C_{Y,S} := \{Z \in \Omega \mid Z \cap B_S(0) = Y \cap B_S(0)\},$$

where $Y \in \Omega$ and $S > 0$ are such that $Y \cap B_S(0) \neq \emptyset$. The next lemma is well known.

Proposition 8. *Every cylinder in Ω is a Cantor set. Moreover, a basis for the topology of Ω is given by sets of the form*

$$\{Z - v \mid Z \in C_{Y,S}, v \in B_\varepsilon(0)\}.$$

In particular, the canonical transversal Ω^0 is a Cantor set.

A *local transversal* in Ω is a clopen (both closed and open) subset of some cylinder in Ω. By Proposition 8, a local transversal C is a Cantor set. This implies that the *recognition radius* defined as

$$\mathrm{rec}(C) := \inf\{S > 0 \mid C_{Y,S} \subseteq C \text{ for all } Y \in C\}$$

is finite. The motivation to define $\mathrm{rec}(C)$ is the following: suppose that a Delone set $Y \in \Omega$ is given and we want to check if Y belongs to C. Then it suffices to look whether the patch $Y \cap \overline{B}_{\mathrm{rec}(C)}(0)$ is equivalent to $Y_i \cap \overline{B}_{\mathrm{rec}(C)}(0)$ for some Y_i. Of course, if $C = C_{Y,S}$, then its recognition radius is smaller than S. Proposition 8 implies also that the collection

$$\{C_{Y,S} \mid Y \in C, S > \mathrm{rec}(C)\}$$

forms a basis for its topology. Indeed, since C is a Cantor set, it is easy to find a finite set $\{Y_1, \ldots, Y_m\}$ in C such that

$$C = \bigcup_{i=1}^{m} C_{Y_i, \mathrm{rec}(C)}.$$

Given a local transversal C and $D \subseteq \mathbb{R}^d$, the following notation will be used throughout the paper:

$$C[D] = \{Y - x \mid Y \in C, x \in D\}.$$

As we define a return vector to a patch, one can define the set of return vectors to a local transversal. Given a local transversal C and a Delone set $Y \in \Omega$, we define

$$\mathcal{R}_C(Y) = \{x \in \mathbb{R}^d \mid Y - x \in C\}.$$

When Y belongs to C, we refer to $\mathcal{R}_C(Y)$ as the *set of return vectors* of Y to C. The following lemma is standard (see, e.g., [C])

Lemma 9. *Let C be a local transversal. Then for each $Y \in C$, the set of return vectors $\mathcal{R}_C(Y)$ is a repetitive Delone set. Moreover, the following quantities*

$$r(C) = \frac{1}{2}\inf\{\|x-y\| \; x, y \in \mathcal{R}_C(Y),\, x \neq y\}, \qquad \textit{and} \tag{3.1}$$

$$R(C) = \inf\{R > 0 \; \mathcal{R}_C(Y) \cap \overline{B}_R(y) \neq \emptyset \textit{ for all } y \in \mathbb{R}^d\}, \tag{3.2}$$

do not depend on the choice of Y in C.

3.1.2. Solenoids and boxes. In this section, we recall some definitions and results of [BBG, BG] that will be used throughout the paper. The hull Ω is locally homeomorphic to the product of a Cantor set and $\mathbb{R}^d$. Moreover, there exists an open cover $\{U_i\}_{i=1}^n$ of Ω such that for each $i \in \{1, \ldots, n\}$, there are $Y_i \in \Omega$, $S_i > 0$ and open sets $D_i \subseteq \mathbb{R}^d$ such that $U_i = C_{Y_i,S_i}[D_i]$ and the map $h_i : D_i \times C_i \to U_i$ defined by $h_i(t, Z) = Z - t$ is a homeomorphism. Furthermore, there are vectors $v_{i,j} \in \mathbb{R}^d$ (depending *only* on i and j) such that the transition maps $h_i^{-1} \circ h_j$ satisfy

$$h_i^{-1} \circ h_j(t, Z) = (t - v_{i,j}, Z - v_{i,j}) \tag{3.3}$$

at all points (t, Z) where the composition is defined. Following [BG], we call such a cover a *$\mathbb{R}^d$-solenoid's atlas*. It induces, among others structures, a laminated structure as follows. First, *slices* are defined as sets of the form $h_i(D_i \times \{Z\})$. Equation (3.3) implies that slices are mapped onto slices. Thus, the *leaves* of Ω are defined as the smallest connected subsets that contain all the slices they intersect. It is not difficult to check, using (3.3), that the leaves coincide with the orbits of Ω.

A *box* in Ω is a set of the form $B := C[D]$ where C is a local transversal in Ω, and $D \subseteq \mathbb{R}^d$ is an open set such that the map from $D \times C$ to B given by $(x, Y) \mapsto Y - x$ is a homeomorphism. This is true, for instance, if $D \subseteq B_{r(C)}(0)$ (*cf.* (3.1)).

3.2. Tower systems

In this section we review the concepts of box decompositions and tower systems introduced in [BBG, BG]. We focus on linearly repetitive Delone sets. The main results of this section can be found in [AC]. For all this section, Ω denotes the hull of an aperiodic repetitive Delone set X.

3.2.1. Box decompositions and derived tilings. A *box decomposition* is a finite and pairwise-disjoint collection of boxes $\mathcal{B} = \{B_1, \ldots, B_t\}$ in Ω such that the closures of the boxes in $\mathcal{B}$ cover the hull. For simplicity, we always write $B_i = C_i[D_i]$, where C_i and D_i are fixed and C_i is contained in B_i. In particular, the set D_i contains 0. We refer to C_i as the *base* of B_i. In this way, we call the union of all C_i the *base* of $\mathcal{B}$. The reasoning for fixing a local transversal in each B_i comes from the fact that box decompositions can be constructed in a canonical way starting from the set $\mathcal{R}_C(Y)$ of return vectors to a given local transversal C [BBG].

An alternative way of understanding a box decomposition is given by a family of tilings, known as *derived tilings*, which are constructed by intersecting the box decomposition with the orbit of each Delone set in the hull.

Let us start by recalling some basic definitions about tilings. A *tile* T in $\mathbb{R}^d$ is a compact set that is the closure of its interior (not necessarily connected). A *tiling* $\mathcal{T}$ of $\mathbb{R}^d$ is a countable collection of tiles that cover $\mathbb{R}^d$ and have pairwise disjoint interiors. Tiles can be *decorated*: they may have a color and/or be punctured at an interior point. Formally, this means that decorated tiles are tuples (T, i, x), where T is a tile, i lies in a finite set of *colors*, and x belongs to the interior of T. Two tiles have the same type if they differ by a translation. If the tiles are punctured, then the translation must also send one puncture to the other, and when they are colored, they must have the same color.

To construct a derived tiling, the idea is to read the intersection of the boxes in the box decomposition with the orbit of a fixed Delone set in the hull. In the sequel, it will be convenient to make the following construction. Let $\{C_i\}_{i=1}^t$ be a collection of local transversals and $\{D_i\}_{i=1}^t$ be a collection of bounded open subsets of $\mathbb{R}^d$ containing 0. Define $\mathcal{B} = \{C_i[D_i]\}_{i=1}^t$ and observe that the sets in $\mathcal{B}$ are not necessarily boxes of Ω. For each $Y \in \Omega$, define the (decorated) *derived collection* of $\mathcal{B}$ at Y by

$$\mathcal{T}_{\mathcal{B}}(Y) := \{(\overline{D_i} + v, i, v) \mid i \in \{1, \dots, t\}, v \in \mathcal{R}_{C_i}(Y)\}.$$

The following lemma gives the relation between box decomposition and tilings.

Lemma 10 ([AC, Lem. 3.1]). *Let $\mathcal{B} = \{C_i[D_i]\}_{i=1}^t$, where the C_i's are local transversals and the D_i's are open bounded subsets of $\mathbb{R}^d$ that contain 0. Then, $\mathcal{B}$ is a box decomposition if and only if $\mathcal{T}_{\mathcal{B}}(Y)$ is a tiling of $\mathbb{R}^d$ for every $Y \in \Omega$. In this case, we call $\mathcal{T}_{\mathcal{B}}(Y)$ the* derived tiling *of $\mathcal{B}$ at Y.*

Proof. It is easy to see that if $\mathcal{B}$ is a box decomposition, then $\mathcal{T}_{\mathcal{B}}(Y)$ is a tiling for every $Y \in \Omega$. We now show the converse. For convenience, set $C = \cup_i C_i$. Fix $Y \in \Omega$ and suppose there are $i, j \in \{1, \dots, t\}$, $Y_1 \in C_i, Y_2 \in C_j$, $x_1 \in D_i$ and $x_2 \in D_j$ such that $Y = Y_1 - x_1 = Y_2 - x_2$. This implies that the tiles $\overline{D_i} - x_1$ and $\overline{D_j} - x_2$ of $\mathcal{T}_{\mathcal{B}}(Y)$ meet an interior point. Since $\mathcal{T}_{\mathcal{B}}(Y)$ is a tiling, these tiles must coincide, and hence $i = j$ and $x_1 = x_2$. We conclude i the maps $h_i : C_i \times D_i \to C_i[D_i]$ given by $(Y, t) \mapsto Y - t$ are one-to-one, and moreover their image are pairwise disjoint. It is then straightforward to check that the maps h_i are homeomorphims. □

3.2.2. Properly nested box decompositions. A box decomposition $\mathcal{B}' = \{C_i'[D_i']\}_{i=1}^{t'}$ is *zoomed out* of another box decomposition $\mathcal{B} = \{C_j[D_j]\}_{j=1}^t$ if the following properties are satisfied:

(Z.1) If $Y \in C_i'$ is such that $Y - x \in C_j - y$ for some $x \in \overline{D_i'}$ and $y \in \overline{D_j}$, then $C_i' - x \subseteq C_j - y$.

(Z.2) If $x \in \partial D_i'$, then there exist j and $y \in \partial D_j$ such that $C_i' - x \subseteq C_j - y$.

(Z.3) For every box B' in $\mathcal{B}'$, there is a box B in $\mathcal{B}$ such that $B \cap B' \neq \emptyset$ and $\partial B \cap \partial B' = \emptyset$.

For each $i \in \{1, \dots, t'\}$ and $j \in \{1, \dots, t\}$ define

$$O_{i,j} = \{x \in D'_i \mid C'_i - x \subseteq C_j\}. \tag{3.4}$$

(Z.4) For each $i \in \{1, \dots, t'\}$ and $j \in \{1, \dots, t\}$,

$$\overline{D'_i} = \bigcup_{j=1}^{t} \bigcup_{x \in O_{i,j}} \overline{D_j} + x,$$

where all the sets in the right-hand side of the equation have pairwise disjoint interiors.

Observe that in the case that D_j is connected, then properties (Z.1) and (Z.2) imply (Z.4).

Since we are considering the C'_i's and C_j's as the bases of the boxes, we ask the following additional property to be satisfied:

(Z.5) The base of $\mathcal{B}'$ is included in the base of $\mathcal{B}$, that is, $\cup_i C'_i \subseteq \cup_j C_j$.

By (Z.4), we have that the tiling $\mathcal{T}_{\mathcal{B}'}(Y)$ is a super-tiling of $\mathcal{T}_{\mathcal{B}}(Y)$ in the sense that each tile T in $\mathcal{T}_{\mathcal{B}'}(Y)$ can be decomposed into a finite set of tiles of $\mathcal{T}_{\mathcal{B}}(Y)$. By (Z.3), one of these tiles is included in the interior of T.

Lemma 11. *For every $j \in \{1, \dots, t\}$ we have*

$$C_j = \bigcup_{i=1}^{t'} \bigcup_{x \in O_{i,j}} C'_i - x.$$

Proof. By the definition of $O_{i,j}$ and (Z.1), it suffices to show that every $Y \in C_j$ belongs to the interior of some box $C'_i[D'_i]$. Suppose not, then $Y \in C'_i - x$ with $x \in \partial D'_i$ for some i since $\mathcal{B}'$ is a box decomposition. Moreover, by (Z.2) we deduce that Y must be in the boundary of some box $B_{j'}$ in $\mathcal{B}$, which gives a contradiction. □

3.3. Tower systems of linearly repetitive Delone system

A *tower system* is a sequence of box decompositions $(\mathcal{B}_n)_{n \in \mathbb{N}}$ such that $\mathcal{B}_{n+1}$ is zoomed out of $\mathcal{B}_n$ for all $n \in \mathbb{N}$. An iteration of the construction of zoomed out box decomposition gives the following result.

Theorem 12 ([BBG]). *The hull of any aperiodic minimal Delone set possesses a tower system.*

We have explained in Section 3.2.1 how to construct a box decomposition and in Section 3.2.2 the notion of zoomed out box decomposition. In this section, we specify the construction of a tower system to the linear repetitive case.

Fot a decreasing sequence $(C_n)_{n \in \mathbb{N}}$ of local transversals with diameter going to 0, and a tower system $(\mathcal{B}_n)_n$, we say that $(\mathcal{B}_n)_n$ is *adapted* to $(C_n)_n$, if for any $n \in \mathbb{N}$ we have $\mathcal{B}_n = \{C_{n,i}[D_{n,i}]\}_{i=1}^{t_n}$ such that $C_n = \cup_i C_{n,i}$ and t_n is a positive integer. In this case, for each $n \in \mathbb{N}^*$ we define, as in (3.4),

$$O^{(n)}_{i,j} = \{x \in D_{n,i} \mid C_{n,i} - x \subseteq C_{n-1,j}\} \tag{3.5}$$

and

$$m_{i,j}^{(n)} = \sharp O_{i,j}^{(n)}$$

for every $i \in \{1, \dots, t_n\}$ and $j \in \{1, \dots, t_{n-1}\}$. The *transition matrix* of level n (associated to the tower system $(\mathcal{B}_n)_{n\in\mathbb{N}}$) is then defined as the matrix $M_n = (m_{i,j}^{(n)})_{i,j}$, so M_n has size $t_n \times t_{n-1}$. From (Z.4), we get

$$\operatorname{vol}(D_{n,i}) = \sum_{j=1}^{t_{n-1}} m_{i,j}^{(n)} \operatorname{vol}(D_{n-1,j}). \tag{3.6}$$

Given a box decomposition $\mathcal{B} = \{C_i[D_i]\}_{i=1}^{t}$, define its external and internal radius by

$$R_{\text{ext}}(\mathcal{B}) = \max_{i\in\{1,\dots,t\}} \inf\{R > 0 : B_R(0) \supseteq D_i\};$$
$$r_{\text{int}}(\mathcal{B}) = \min_{i\in\{1,\dots,t\}} \sup\{r > 0 : B_r(0) \subseteq D_i\},$$

respectively. Define also $\operatorname{rec}(\mathcal{B}) = \max_{i\in\{1,\dots,t\}} \operatorname{rec}(C_i)$.

With all theses definitions, we can state the following result for aperiodic linearly repetitive Delone systems.

Theorem 13 ([AC, Thm. 3.4]). *Let X be an aperiodic linearly repetitive Delone set with constant $L > 1$ and $0 \in X$. Given $K \geq 6L(L+1)^2$ and $s_0 > 0$, set $s_n := K^n s_0$ and $C_n := C_{X,s_n}$ for all $n \in \mathbb{N}$. Then, there exists a tower system $(\mathcal{B}_n)_n$ of Ω adapted to $(C_n)_{n\in\mathbb{N}}$ that satisfies the following additional properties:*

i) *for every $n \geq 0$, $C_{n+1} \subseteq C_{n,1}$;*

ii) *there exist constants*

$$K_1 := \frac{1}{2(L+1)} - \frac{L}{K-1} \quad \text{and} \quad K_2 := \frac{LK}{K-1},$$

which satisfy $0 < K_1 < 1 < K_2$, such that for every $n \in \mathbb{N}$ we have

$$K_1 s_n \leq r_{\text{int}}(\mathcal{B}_n) < R_{\text{ext}}(\mathcal{B}_n) \leq K_2 s_n; \tag{3.7}$$

iii) *for every $n \in \mathbb{N}$,*

$$\operatorname{rec}(\mathcal{B}_n) \leq (2L+1)s_n. \tag{3.8}$$

As an application of this result, we have the nice following structure.

Theorem 14. *Let X be an aperiodic linearly repetitive Delone set. Then, the tower system of Ω obtained in Theorem 13 satisfies the following:*

1. *For every $n \in \mathbb{N}^*$, the matrix $M_n = (m_{i,j}^{(n)})_{i,j}$ has strictly positive coefficients;*
2. *The matrices $\{M_n\}_{n\in\mathbb{N}^*}$ are uniformly bounded in size and norm.*

In the self-similar case, the family of matrices $\{M_n\}$ can be reduced to only one element.

Proof. Take the notations of Theorem 13. Indeed, by the definition of linearl repetitivity, we have $M_X(\mathrm{rec}(\mathcal{B}_n)) \leq L\,\mathrm{rec}(\mathcal{B}_n)$ for all $n \in \mathbb{N}^*$. Combining this with (3.8), the left-hand inequality of (3.7) and the definition of s_n we get

$$M_X(\mathrm{rec}(\mathcal{B}_n)) \leq \frac{L(2L+1)}{KK_1} r_{\mathrm{int}}(\mathcal{B}_{n+1}).$$

Since $K \geq 6L(L+1)^2$, it follows that $L(2L+1) \leq K_1 K$ and we obtain for all $n \geq 0$

$$M_X(\mathrm{rec}(\mathcal{B}_n)) \leq r_{\mathrm{int}}(\mathcal{B}_{n+1}).$$

Thus any $\mathrm{rec}(\mathcal{B}_n)$-patch occurs in a set $D_{n+1,i} \cap Y$ for any $Y \in C_{n+1,i}$, and the coefficients $m_{i,j}^{(n)}$ are positive. Moreover, since $D_{n,i}$ is included in a ball of radius $R_{\mathrm{ext}}(\mathcal{B}_{n-1})$ and each $D_{n-1,j}$ contains a ball of radius $r_{\mathrm{int}}(\mathcal{B}_{n-1})$, we deduce from (3.6) that

$$\sum_{j=1}^{t_{n-1}} m_{i,j}^{(n)} \leq \left(\frac{R_{\mathrm{ext}}(\mathcal{B}_{n-1})}{r_{\mathrm{int}}(\mathcal{B}_{n-1})}\right)^d \leq \left(K\frac{K_2}{K_1}\right)^d.$$

So we get that the matrices $\{M_n\}_n$ are uniformly bounded. □

4. Ergodic properties of linearly repetitive system

4.1. Background on transverse invariant measure

A Borel measure μ on the hull Ω of a repetitive Delone set is *translation invariant* if $\mu(B - v) = \mu(B)$ for every Borel set B and $v \in \mathbb{R}^d$. It is well known that any continuous $\mathbb{R}^d$ action on a compact space admits an invariant measure.

Let C be a local transversal and $0 < r < r(C)$. Each translation invariant measure μ induces a measure ν on C (see [Gh] for the general construction): given a Borel subset V of C, its *transverse measure* is defined by

$$\nu(V) = \frac{\mu(V[B_r(0)])}{\mathrm{vol}(B_r(0))},$$

where vol denotes the Euclidean volume in $\mathbb{R}^d$. This gives a measure on each C, which does not depend on small r. The collection of all measures defined in this way is called the *transverse invariant measure* induced by μ. It is invariant in the sense that if V is a Borel subset of C and $x \in \mathbb{R}^d$ is such that $V - x$ is a Borel subset of another local transversal C', then $\nu(V - x) = \nu(V)$. Conversely, the measure μ of any box B written as $C[D]$ may be computed by the equation

$$\mu(C[D]) = \mathrm{vol}(D) \times \nu(C).$$

For a tower system $(\mathcal{B}_n)_{n\geq 0}$ where $\mathcal{B}_n = \{C_{n,i}[D_{n,i}]\}_{i=1}^{t_n}$ from (Z.4), Lemma 11 and the definition of transverse invariant measures, we get

$$\nu(C_{n-1,j}) = \sum_{i=1}^{t_n} \nu(C_{n,i}) m_{i,j}^{(n)}. \tag{4.1}$$

Fix $n \in \mathbb{N}$. From the relation $\mu(C_{n,i}[D_{n,i}]) = \mathrm{vol}(D_{n,i})\nu(C_{n,i})$ and the fact that $\mathcal{B}_n$ is a box decomposition, it follows that

$$\sum_{j=1}^{t_n} \mathrm{vol}(D_{n,j})\nu(C_{n,j}) = 1. \tag{4.2}$$

4.2. Unique ergodicity and speed of convergence

When the system $(\Omega, \mathbb{R}^d)$ has a unique translation invariant probability measure, the system is called *uniquely ergodic*. The unique ergodicity implies combinatorial properties for the Delone set. The dynamical system $(\Omega, \mathbb{R}^d)$ is uniquely ergodic, if and only if any Delone set $X \in \Omega$ has *uniform patch frequencies*, i.e., any patch $\mathtt{P}$ occurs with a positive frequency; more precisely: Let $X_{\mathtt{P}}$ be the set of occurrences of the patch $\mathtt{P}$ in X, and let $(D_N)_N$ be a nested sequence of d-cube D_N of side N, then the following limit exists.

$$\lim_{N\to\infty} \frac{\sharp X_{\mathtt{P}} \cap D_N}{\mathrm{vol}(D_N)} =: \mathrm{freq}(\mathtt{P}).$$

The number $\mathrm{freq}(\mathtt{P})$ is called the *frequency* of $\mathtt{P}$. Notice the difference with the standard Birkhoff ergodic Theorem that asserts a convergence only for almost all Delone set of the hull.

Theorem 15. *Let X be an aperiodic linearly repetitive Delone set of $\mathbb{R}^d$ and Ω its hull. Then the system $(\Omega, \mathbb{R}^d)$ is uniquely ergodic.*

The original proof is due to Lagarias and Pleasants in [LP2]. By using the identification between a transverse invariant measure and the inverse limit of top homologies of branched manifolds, the authors in [BBG] show that in the case described in Theorem 14, the system is uniquely ergodic. This proof is independent of the original one.

Actually for linearly repetitive system, we can be much more precise and give informations on the speed of convergence of the limit. For instance the following is a stronger result of Lagarias and Pleasants [LP2], that implies the unique ergodicity.

Theorem 16 ([LP2]). *Let X be a linearly repetitive Delone set of $\mathbb{R}^d$. There exists a $\delta > 0$ such that, for every patch $\mathtt{P}$ of X, there is a number* $\mathrm{freq}(\mathtt{P})$ *so that*

$$\left| \frac{X_{\mathtt{P}} \cap Dom_N}{\mathrm{vol}(Dom_N)} - \mathrm{freq}(\mathtt{P}) \right| = O(N^{-\delta}),$$

where Dom_N is either a d-cube with side N or a ball of radius N. The O-constant may depend on the patch $\mathtt{P}$.

In [AC], a proof of this theorem is given using the structure Theorem 14 and relating the constant δ with the matrices M_n by the following way

$$\delta = d - \log_K \left(1 - \sup_n ||M_n||_1^{-1} ||M_{n+1}||_1^{-1} \right),$$

where $\log_K$ denotes the logarithm in base K.

4.3. Non-mixing properties

A translation invariant probability measure μ on a the hull Ω of a Delone set is said to be *measurably strongly mixing* if for any Borel sets A, B in Ω,

$$\lim_{\|v\|\to\infty} \mu((A-v)\cap B) = \mu(A)\mu(B). \tag{4.3}$$

In this section, we show the following proposition which is analogous to theorem of Dekking and Keane [DK] for substitutive subshifts.

Proposition 17 ([C0]). *Let X be a linearly repetitive Delone set of $\mathbb{R}^d$ and Ω its hull. Then the system $(\Omega, \mathbb{R}^d)$ is not measurably strongly mixing.*

The proof's strategy is the same as for self-similar tiling in [So1] or for linear recurrent Cantor system in [CDHM]. But we need sharp estimates on the transverse measures of clopen sets, provided by Theorem 13.

Assume that the Delone set X is aperiodic and linearly repetitive with constant L. Let μ be the unique translation invariant probability measure on the hull Ω, and let ν be the associated transverse invariant measure. Let $(\mathcal{B}_n)_{n\geq 0}$ be the tower system given by Theorem 13 where for each integer n, $\mathcal{B}_n = \{C_{n,i}[D_{n,i}]\}_{i=1}^{t_n}$.

Lemma 18. *For the tower system of Ω given by Theorem* 13*, we have*

$$\inf_{\substack{n\geq 1\\ 1\leq i\leq n}} \mathrm{vol}(D_{n,i})\nu(C_{n,i}) > \left(\frac{K_1}{KK_2}\right)^d =: c > 0.$$

Proof. With equation (4.1), for any $1 \leq i \leq t_n$, we get

$$\nu(C_{n,i}) \geq \sum_{j=1}^{t_{n+1}} \nu(C_{n+1,j}). \tag{4.4}$$

By definition, for any $1 \leq i \leq t_n$, the domain $D_{n,i}$ contains a ball or radius $r_{\mathrm{int}}(\mathcal{B}_n)$ and for $1 \leq j \leq t_{n+1}$ the domain $D_{n+1,j}$ is included in a ball of radius $R_{\mathrm{ext}}(\mathcal{B}_{n+1})$. Thus, as in the proof of Theorem 14, we deduce from Theorem 13

$$\frac{\mathrm{vol}(D_{n+1,j})}{\mathrm{vol}(D_{n,i})} \leq \left(\frac{R_{\mathrm{ext}}(\mathcal{B}_{n-1})}{r_{\mathrm{int}}(\mathcal{B}_{n-1})}\right)^d \leq \left(K\frac{K_2}{K_1}\right)^d = c^{-1}. \tag{4.5}$$

Thus it follows from (4.2), that for any $n \geq 0$ and $1 \leq i \leq t_n$

$$\mathrm{vol}(D_{n,i})\nu(C_{n,i}) \geq \sum_{j=1}^{t_{n+1}} c\mathrm{vol}(D_{n+1,j})\nu(C_{n+1,j}) = c. \qquad \square$$

For the tower system $(\mathcal{B}_n)_n$, we define as in Definition 3.4, for integers $p \geq n > 0$

$$O_{i,j}^{(p,n)} := \{x \in D_{p,i} \mid C_{p,i} - x \subseteq C_{n-1,j}\}, \text{ for } 1 \leq i \leq t_p; 1 \leq j \leq t_{n-1} \tag{4.6}$$

and

$$m_{i,j}^{(p,n)} = \sharp O_{i,j}^{(p,n)}.$$

Then it is straightforward to check that the $t_p \times t_{n-1}$ matrix satisfies

$$(m_{i,j}^{(p,n)})_{i,j} = M_p \cdots M_n.$$

Lemma 19. *For the tower system of Ω given by Theorem 13, we have for $n \geq 2$, and $1 \leq j \leq t_n$*

$$\liminf_{p\to+\infty} \min_{1\leq i\leq t_p} \frac{m_{i,j}^{(p,n)}}{\mathrm{vol}(D_{i,p})} \geq \nu(C_{n-1,j}) \left(\frac{K_1}{K_2}\right)^d c.$$

Proof. Let $X \in \cap_{n\geq 0} C(n)$. By the unique ergodicity, we have

$$\lim_{R\to+\infty} \frac{1}{\mathrm{vol}(B_R(0))} \sharp\{B_R(0) \cap \mathcal{R}_{C_{n,j}}(X)\} = \nu(C_{n,j}). \tag{4.7}$$

Since for every $p > n$, the set $C_p \subset C_{p-1,1}$, we get for any $1 \leq i \leq t_p$,

$$m_{i,j}^{(p,n)} \geq m_{1,j}^{(p-1,n)} \geq \sharp\{D_{p-1,j} \cap \mathcal{R}_{C_{n-1,j}}(X)\}.$$

Hence we conclude by this inequality and inequality (4.5) that

$$\begin{aligned}
\liminf_{p\to+\infty} \min_{1\leq i\leq t_p} \frac{m_{i,j}^{(p,n)}}{\mathrm{vol}(D_{p,i})} &\geq \liminf_{p\to+\infty} \frac{\sharp\{D_{p-1,j} \cap \mathcal{R}_{C_{n-1,j}}(X)\}}{\mathrm{vol}(D_{p,i})} \\
&\geq c \liminf_{p} \frac{\sharp\{D_{p-1,j} \cap \mathcal{R}_{C_{n-1,j}}(X)\}}{\mathrm{vol}(D_{p-1,j})} \\
&\geq c \lim_{p} \frac{\sharp\{B_{r_{\mathrm{int}}(\mathcal{B}_{p-1})}(0) \cap \mathcal{R}_{C_{n-1,j}}(X)\}}{\mathrm{vol}(B_{\frac{K_2}{K_1} r_{\mathrm{int}}(\mathcal{B}_{p-1})}(0))},
\end{aligned}$$

since $D_{p-1,j}$ contains the ball $B_{r_{\mathrm{int}}(\mathcal{B}_{p-1})}(0)$ and is contained in the ball

$$B_{R_{\mathrm{ext}}(\mathcal{B}_{p-1})}(0) \subset B_{\frac{K_2}{K_1} r_{\mathrm{int}}(\mathcal{B}_{p-1})}(0).$$

We obtain the conclusion by the equality (4.7). □

Now we are able to prove Proposition 17.

Proof of Proposition 17. Let n be an integer such that $\nu(C_n) < \left(\frac{K_1}{K_2}\right)^d c^2$. For $p \geq n$, Let $\mathcal{F}_{p,1} \subset \mathbb{R}^d$ be the set of vector v such that there exists a $1 \leq j \leq t_p$ satisfying $C_{p,1} - v \cap C_{p,j} \neq \emptyset$ and $D_{p,j} - v \cap D_{p,1} \neq \emptyset$. Let $\tilde{C}(n, v) = (C_{n,1} - v) \cap C_{n,1}$. We will show that

$$\liminf_{p\to\infty} \inf_{v\in\mathcal{F}_{p,1}} \nu(\tilde{C}(n, v)) > \nu(C_{n,1})^2$$

which implies that the system $(\Omega, \mathbb{R}^d)$ is not strongly mixing.
For $x \in O_{1,1}^{(p,n+1)} = \{x \in D_{p,1} \mid C_{p,1} - x \subseteq C_{n,1}\}$, and $v \in \mathcal{F}_{p,1}$, we have by (Z.1) and by i) in Theorem 13

$$C_{p+1,1} - (v + x) \subset C_{p+1} - x \subset C_{p,1} - x \subset C_{n,1}.$$

Thus for any $x \in O_{1,1}^{(p,n+1)}$ and $v \in \mathcal{F}_{p,1}$ we get $C_{p+1,1} - x \subset \tilde{C}(n,v)$. Then

$$\nu(\tilde{C}(n,v)) \geq \sharp O_{1,1}^{(p,n+1)} \nu(C_{p+1,1}) = m_{1,1}^{(p,n+1)} \nu(C_{p+1,1}).$$

By Lemma 19, we obtain

$$\begin{aligned}
&\liminf_{p\to\infty} \inf_{v\in\mathcal{F}_{p,1}} \nu(\tilde{C}(n,v)) \\
&\geq \liminf_{p\to\infty} \frac{m_{1,1}^{(p,n+1)}}{\operatorname{vol}(D_{1,p})} \nu(C_{p+1,1}) \operatorname{vol}(D_{1,p}) \\
&\geq \nu(C_{n,1}) c \left(\frac{K_1}{K_2}\right)^d \liminf_{p\to\infty} \nu(C_{p+1,1}) \operatorname{vol}(D_{1,p}) \\
&\geq \nu(C_{n,1}) c \left(\frac{K_1}{K_2}\right)^d \liminf_{p\to\infty} \nu(C_{p+1,1}) \operatorname{vol}(D_{1,p+1}) c \text{ by inequality (4.5)} \\
&\geq \nu(C_{n,1}) \left(\frac{K_1}{K_2}\right)^d c^2 > \nu(C_{n,1})^2.
\end{aligned}$$

□

4.4. Subadditive ergodic theorem

In Section 4.2 we recall that the linearly repetitive systems are uniquely ergodic. Actually such systems satisfy also a subadditive ergodic theorem. Let $\mathcal{B}(\mathbb{R}^d)$ denotes the family of bounded subsets in $\mathbb{R}^d$. A real-valued function $F\colon \mathcal{B}(\mathbb{R}^d) \to \mathbb{R}$ is called *subadditive* if

$$F(Q_1 \cup Q_2) \leq F(Q_1) + F(Q_2)$$

for any disjoint sets $Q_1, Q_2 \in \mathcal{B}(\mathbb{R}^d)$. For a Delone set X, the function F is called X-*invariant* if

$$F(Q) = F(Q+t) \quad \text{whenever } Q \in \mathcal{B}(\mathbb{R}^d) \text{ and } t + (Q \cap X) = (t+Q) \cap X.$$

For instance, given a patch $\mathtt{P}$ of the Delone set X, the function $B \in \mathcal{B}(\mathbb{R}^d) \mapsto -\sharp X_{\mathtt{P}} \cap B$ where $X_{\mathtt{P}}$ denotes the set of occurrences of the patches $\mathtt{P}$ in X, is a subadditive X-invariant function.

Theorem 20 ([DL, BBL]). *Let X be a linearly repetitive Delone set in $\mathbb{R}^d$. Then X satisfies the* uniform ergodic theorem*:* i.e.*, for any X-invariant subadditive function F and any nested sequence $(D_n)_n$ of d-cubes with side-lengths going to infinity as n goes to infinity, the following limit exists*

$$\lim_{n\to+\infty} \frac{F(D_n)}{\operatorname{vol}(D_n)},$$

and is independent of the sequence $(D_n)_n$.

It is then easy to deduce from this result that the associated dynamical system is uniquely ergodic. The converse is false, in [DL], the authors give an example of a Sturmian sequence that does not satisfy the subadditive ergodic theorem. They prove also a more stronger form of this theorem.

The *lower density* $\underline{\nu}(\mathtt{P})$ of a R-patch $\mathtt{P}$ is the quantity

$$\underline{\nu}(\mathtt{P}) := \lim\inf_{n\to\infty} \frac{\sharp X_{\mathtt{P}} \cap B_n(0)}{\mathrm{vol}(B_n(0))}\,\mathrm{vol}(B_R(0)).$$

The results in [BBL] have this direct corollary.

Proposition 21. *If X is a repetitive (r_X, R_X) Delone set verifying the uniform subadditive ergodic theorem, then X satisfies* positivity of weights, i.e.,

$$\inf_{P \text{ is an } R\text{-patch, } R \geq R_X} \underline{\nu}(P) > 0.$$

Notice that in dimension 1, the positivity of weights property is sufficient to ensure the unique ergodiciy (see [Bo]). Actually, one can deduce from Lemma 18 that a linearly repetitive Delone set satisfies the positivity of weights.

4.5. A characterization of linear repetitivity

In [Le02], D. Lenz characterizes the subshifts that admit a uniform subadditive ergodic Theorem by uniform positivity of weights. This can be considered as an averaged version of linear repetitivity. For Delone systems, it is shown in [Bes, BBL] that the linear repetitivity is equivalent to positivity of weights plus some balancedness of the shape of patterns. For a Voronoï cell V of a Delone set, let us define:

$$r_{\mathrm{int}} := \sup\{r > 0; V \text{ contains a ball of radius } r\}.$$
$$R_{\mathrm{ext}} := \inf\{R > 0; V \text{ is contained in a ball of radius } R\}.$$

The *distorsion* of V is the constant $\lambda(V) := R_{\mathrm{ext}}(V)/r_{\mathrm{int}}(V)$.

Theorem 22 ([BBL]). *Let X be an aperiodic Delone set in $\mathbb{R}^d$ of finite type. Then X is linearly repetitive if and only if for any R-patch P of X, $R > 0$: the set X_P of occurrences of P is a (r_P, R_P)-Delone set such that*

(i) $\sup_{P, x\in X_P} \lambda(V_x) < +\infty$ *where V_x denotes the Voronoï cell of x.*
(ii) *The Delone set X satisfies the positivity of weights* (*see Proposition* 21).

One can find in [BBL] another similar equivalent condition to linear repetitivity. Notice that in dimension $d = 1$, the distorsion of any compact Voronoï cell is equal to 1. Thus the condition (ii) is equivalent to the linear repetitivity.

For an aperiodic linearly repetitive Delone set, the properties (i)–(ii) arise from the properties recalled in Subsections 2.1 and 4.4.

Let us also mention in Chapter **On the non commutative geometry for tilings**, a characterization of Sturmian sequences that are linearly repetitive by using metrics arising from the Connes distance.

5. Factors of linearly repetitive system

A *factor map* between two Delone systems $(\Omega_1, \mathbb{R}^d)$ and $(\Omega_2, \mathbb{R}^d)$ is a continuous surjective map $\pi : \Omega_1 \to \Omega_2$ such that $\pi(X - v) = \pi(X) - v$, for every $X \in \Omega_1$ and $v \in \mathbb{R}^d$.

In symbolic dynamics it is well known that topological factor maps between subshifts are always given by sliding-block-codes. An equivalent notion for the Delone system is the *local derivability*: *i.e.*, there exists a constant $s_0 > 0$ such that for any radius $R > 0$, if two Delone sets $X, Y \in \Omega_1$ satisfy $X \cap B_{R+s_0}(0) = Y \cap B_{R+s_0}(0)$ then $\pi(X) \cap B_R(0) = \pi(Y) \cap B_R(0)$. However there are examples of factor maps on Delone systems that are not sliding-block codes ([Pe, RS]). Nevertheless, the following lemma shows that factor maps between Delone systems are not far from being sliding-block-codes. Similar results can be found in [CD, CDP, HRS].

Lemma 23. *Let X_1 and X_2 be two Delone sets. Suppose X_1 has finite local complexity and $\pi : \Omega_{X_1} \to \Omega_{X_2}$ is a factor map. Then, there exists a constant $s_0 > 0$ such that for every $\varepsilon > 0$, there exists $R_\varepsilon > 0$ satisfying the following: For any $R \geq R_\varepsilon$, if X and X' in Ω_{X_1} satisfy*

$$X \cap B_{R+s_0}(0) = X' \cap B_{R+s_0}(0),$$

then

$$(\pi(X) - v) \cap B_R(0) = \pi(X') \cap B_R(0)$$

for some $v \in B_\varepsilon(0)$.

Proof. The Delone set X_2 has also finite local complexity because Ω_{X_2} is compact. Let r_0 and R_0 be positive constants such that X_2 is a (r_0, R_0)-Delone set. Since all the elements of Ω_{X_2} are (r_0, R_0)-Delone sets, if two different points y_1, y_2 of $\mathbb{R}^d$ satisfy $(X - y_1) \cap B_R(a) = (X - y_2) \cap B_R(a)$ for some $X \in \Omega_{X_2}$, $a \in \mathbb{R}^d$ and $R > R_0$, then $\|y_1 - y_2\| \geq \frac{r_0}{2}$ (for the details see [So1]).

Let $0 < \delta_0 < \min\{\frac{r_0}{4}, \frac{1}{R_0}\}$. Since π is uniformly continuous, there exists $s_0 > 1$ such that if X and X' in Ω_{X_1} verify $X \cap B_{s_0}(0) = X' \cap B_{s_0}(0)$ then

$$(\pi(X) - v) \cap B_{\frac{1}{\delta_0}}(0) = \pi(X') \cap B_{\frac{1}{\delta_0}}(0),$$

for some $v \in B_{\delta_0}(0)$. Let $0 < \varepsilon < \delta_0$. By uniform continuity of π, there exists $0 < \delta < \frac{1}{s_0}$ such that if X and X' in Ω_{X_1} verify $X \cap B_{\frac{1}{\delta}}(0) = X' \cap B_{\frac{1}{\delta}}(0)$ then

$$(\pi(X) - v) \cap B_{\frac{1}{\varepsilon}}(0) = \pi(X') \cap B_{\frac{1}{\varepsilon}}(0), \tag{5.1}$$

for some $v \in B_\varepsilon(0)$. Now fix $R \geq R_\varepsilon = \frac{1}{\delta} - s_0$, and let X and X' be two Delone sets in Ω_{X_1} satisfying

$$X \cap B_{R+s_0}(0) = X' \cap B_{R+s_0}(0). \tag{5.2}$$

Observe that X and X' satisfy (5.1), and $(X - a) \cap B_{s_0}(0) = (X' - a) \cap B_{s_0}(0)$, for every a in $B_R(0)$. The choice of s_0 ensures that

$$(\pi(X) - a - t(a)) \cap B_{\frac{1}{\delta_0}}(0) = (\pi(X') - a) \cap B_{\frac{1}{\delta_0}}(0), \tag{5.3}$$

for some $t(a) \in B_{\delta_0}(0)$. Let us prove the map $a \mapsto t(a)$ is locally constant. For $a \in B_R(0)$, let $0 < s_a < \frac{1}{\delta_0} - R_0$ be such that $B_{s_a}(a) \subseteq B_R(0)$. Every $a' \in B_{s_a}(0)$ verifies $B_{\frac{1}{\delta_0}-\|a'\|}(-a') \subset B_{\frac{1}{\delta_0}}(0)$. Let $a' \in B_{s_a}(0)$. This inclusion and (5.3) imply

$$(\pi(X) - a - a' - t(a)) \cap B_{\frac{1}{\delta_0}-\|a'\|}(-a') = (\pi(X') - a - a') \cap B_{\frac{1}{\delta_0}-\|a'\|}(-a'). \quad (5.4)$$

On the other hand, from the definition of the map $a \mapsto t(a)$ we deduce

$$(\pi(X) - a - a' - t(a + a')) \cap B_{\frac{1}{\delta_0}}(0) = (\pi(X') - a - a') \cap B_{\frac{1}{\delta_0}}(0),$$

which implies

$$(\pi(X) - a - a' - t(a + a')) \cap B_{\frac{1}{\delta_0}-\|a'\|}(-a') = (\pi(X') - a - a') \cap B_{\frac{1}{\delta_0}-\|a'\|}(-a'). \quad (5.5)$$

Since $\|t(a) - t(a + a')\| \leq \frac{r_0}{2}$, from equations (5.4), (5.5) and the remark of the beginning of the proof we conclude $t(a) = t(a + a')$ for every $a' \in B_s(0)$. Therefore the map $a \mapsto t(a)$ is constant on $B_{s_a}(a)$.

Furthermore, due to $\delta_0 > \varepsilon$ and (5.2), Equation (5.1) implies there exists $v \in B_\varepsilon(0)$ such that

$$(\pi(X) - v) \cap B_{\frac{1}{\delta_0}}(0) = \pi(X') \cap B_{\frac{1}{\delta_0}}(0). \quad (5.6)$$

For $a = 0$, from (5.3) and (5.6) we have that $t(0) = v$ or $\|v - t(0)\| \geq \frac{r_0}{2}$. Since $\|t(0) - v\| \leq \delta_0 + \varepsilon < 2\delta_0 < \frac{r_0}{2}$, we conclude $t(0) = v$ and then $t(a) = v$ for every $a \in B_R(0)$. This property together with (5.3) and (5.6) imply that

$$(\pi(X) - v) \cap B_R(0) = \pi(X') \cap B_R(0).$$

This conclude the proof. □

Lemma 24 ([CD, Lem. 3]). *Let X_1 and X_2 be two Delone sets with finite local complexity. If $\pi : \Omega_{X_1} \to \Omega_{X_2}$ is a factor map and X_1 is linearly repetitive, then $(\Omega_{X_2}, \mathbb{R}^d)$ is linearly repetitive.*

Proof. Let $X \in \Omega_{X_1}$. Consider $0 < \varepsilon < 1$ and $s_0, R(\varepsilon) > 0$ the positive constants of Lemma 23 associated to ε. Since X is linearly repetitive with some constant L, for any $y \in \mathbb{R}^d$ there exists $v \in B_{L(R+s_0)}(y)$ such that $(X - v) \cap B_{R+s_0}(0) = X \cap B_{R+s_0}(0)$. From Lemma 23, there exists $t \in B_\varepsilon(0)$ such that $(\pi(X) - v - t) \cap B_R(0) = \pi(X) \cap B_R(0)$. This implies that any ball of radius $L(R + s_0) + 2\varepsilon$ in $\pi(X)$ contains a copy of $\pi(X) \cap B_R(0)$. Since $Ls_0 + 2\varepsilon$ is smaller than the constant $Ls_0 + 2$, it follows that $\pi(X)$ is linearly repetitive. □

Actually from the proofs of Lemmas 4 and 24 we can get a uniform bound on the linear repetitivity constant of the factor system.

Lemma 25. *Let X_1 and X_2 be two Delone sets with finite local complexity. If $\pi : \Omega_{X_1} \to \Omega_{X_2}$ is a factor map and X_1 is linearly repetitive with constant $L > 1$, then there exists $R_\pi > 0$ such that for every $R > R_\pi$ and every R-patch $\mathtt{P}$ of X_2, a copy of $\mathtt{P}$ appears in every ball of radius $3LR$ of X_2 and any two occurrences of $\mathtt{P}$ in X_2 are at distance at least $R/4L$.*

5.1. Finite number of aperiodic Delone systems as factors

The aim of this section is to prove the following theorem that is a generalization of a result in [Du1] in the context of subshifts.

Theorem 26 ([CDP, Thm. 12]). *Let $L > 1, d \geq 1$. There exists a constant $N(L,d)$ such that any linearly repetitive Delone set X of $\mathbb{R}^d$ with constant L, has at most $N(L,d)$ aperiodic Delone system factors of $(\Omega_X, \mathbb{R}^d)$ up to conjugacy.*

The bound $N(L,d)$ is essentially due to the constants arising in Lemmas 6 and 7. The proof relies on a generalization of these lemmas and on the specific structure of the factor maps for linearly repetitive Delone systems.

The next result says that factor maps between linearly repetitive Delone systems are finite-to-one. A proof of that result in the context of subshifts and Delone systems can be found in [Du1] and in [CDP, Prop. 5] respectively. Here we include the proof in the case where the factor map is a sliding-block-code.

Proposition 27. *Let X be a linearly repetitive Delone set with constant L. There exists a constant $C > 0$ (depending only on L) such that If X' is an aperiodic Delone set and $\pi : (\Omega_X, \mathbb{R}^d) \to (\Omega_{X'}, \mathbb{R}^d)$ is a factor map, then for every $Y \in \Omega_{X'}$, the fiber $\pi^{-1}(\{Y\})$ contains at most C elements.*

Proof. For simplicity we will assume that π is a sliding-block-code. That means there exists $s_0 > 0$ such that if X_1 and $X_2 \in \Omega_X$ verify $X_1 \cap B_{R+s_0}(0) = X_2 \cap B_{R+s_0}(0)$ for an $R > 0$, then $\pi(X_1) \cap B_R(0) = \pi(X_2) \cap B_R(0)$. From Lemma 24 the Delone set X' is linearly repetitive, and if R is sufficiently large, Lemma 25 implies that for any $x \in \mathbb{R}^d$ a copy of the patch $X' \cap B_R(x)$ appears in $X' \cap B_{3LR}(y)$, for every $y \in \mathbb{R}^d$. Let $Y \in \Omega_{X'}$ and $X_1, \ldots, X_n$ be different Delone sets in $\in \pi^{-1}(\{Y\})$. Because these Delone sets are different, for every sufficiently large R, the patches $X_i \cap B_R(0)$ are pairwise distinct. Linear repetitivity of X ensures the existence of points $v_1, \ldots, v_n \in B_{LR}(0)$ such that each $X - v_i \cap B_R(0)$ is a copy of $X_i \cap B_R(0)$, for every $1 \leq i \leq n$. This implies that $\pi(X) - v_i \cap B_{R-s_0}(0) = Y \cap B_{R-s_0}(0)$. From this and Lemma 25 we get that $\|v_i - v_j\| \geq \frac{R-s_0}{4L}$, from which we deduce that $n \leq C$, where C is a constant that depends only on L. □

The following proposition is a straightforward generalization of Lemma 21 in [Du1]. A proof in our setting can be found in [CDP, Prop. 6]. Here we omit the proof.

Proposition 28. *Let $(\Omega, \mathbb{R}^d)$ be a minimal Delone system and $\phi_1 : (\Omega, \mathbb{R}^d) \to (\Omega_1, \mathbb{R}^d)$, $\phi_2 : (\Omega, \mathbb{R}^d) \to (\Omega_2, \mathbb{R}^d)$ be two factor maps. Suppose that $(\Omega_2, \mathbb{R}^d)$ is non-periodic and ϕ_1 is finite-to-one. If there exist $X, Y \in \Omega$ and $v \in \mathbb{R}^d$ such that $\phi_1(X) = \phi_1(Y)$ and $\phi_2(X) = \phi_2(Y - v)$, then $v = 0$.*

We have already defined the notion of return vector of a patch, now let us define the notion of return vector of a Voronoï cell of a patch. For a patch $\mathtt{P}$ of X and $v \in X_{\mathtt{P}}$, $V_{\mathtt{P},v}$ denotes the Voronoï cell of the point v of the Delone set $X_{\mathtt{P}}$. We

say that $w \in \mathbb{R}^d$ is a *return vector of* $V_{\mathtt{P},v} \cap X$ if $(X-w) \cap V_{\mathtt{P},v} = X \cap V_{\mathtt{P},v}$. We set for $n \geq 1, v \in X_{\mathtt{P}}$,

$$\mathtt{P}_{n,w,v} \text{ the patch } (X - w - v) \cap B_{L^n R}(0).$$

Notice that $\mathtt{P}_{n,w,v} + v + w$ is a patch of X. When there is no confusion about n and v, we write $\mathtt{P}_w$ instead of $\mathtt{P}_{n,w,v}$.

The following lemma generalizes Lemma 6.

Lemma 29. *Let $n \in \mathbb{N}^*$ and X be an aperiodic linearly repetitive Delone set with constant L. There exists a constant $C(n,L) > 0$ such that for every sufficiently large $R > 0$ and every R-patch $\mathtt{P}$, the collection $\{\mathtt{P}_{n,w,v} : w$ is a return vector of $V_{\mathtt{P},v} \cap X\}$ has at most $C(n,L)$ elements, for every $v \in X_{\mathtt{P}}$.*

Proof. Let $\mathtt{P} = X \cap B_R(x_{\mathtt{P}})$ and $v \in X_{\mathtt{P}}$. Lemma 4 implies that the Voronoï cell $V_{\mathtt{P},v}$ contains the ball $B_{\frac{R}{2(L+1)}}(v)$. Then for every pair of return vectors u and w of $V_{\mathtt{P},v}$, the patches $\mathtt{P}_u$ and $\mathtt{P}_w$ coincides at the ball $B_{\frac{R}{2(L+1)}}(0)$. The proof concludes using the fact that in $X \cap B_{2L(L^n R)}(0)$ there is at least one copy of each patch $\mathtt{P}_w$, $\mathtt{P}_u$ and applying Lemma 4 to the return vectors of the patch $\mathtt{P}_w \cap B_{\frac{R}{2(L+1)}}(0)$. □

Proof of Theorem 26. It is enough to suppose that X is an aperiodic linearly repetitive Delone set with constant $L > 1$. Let $n \in \mathbb{N}$ be such that

$$L^n - 1 - 12L - 64L^2 > 1. \tag{5.7}$$

We call $M(n,L)$ the number of coverings of a set with $c(L)c(n,L)$ elements, where $c(L)$ and $c(n,L)$ are the constants of Lemma 6 and Lemma 29 respectively. For every $1 \leq i \leq M(n,L)+1$, let X_i be a non-periodic Delone set such that there exists a topological factor map $\pi_i : \Omega_X \to \Omega_{X_i}$, and let $X_0 = X$. We will show there exist $1 \leq i < j \leq M(n,L)+1$ such that $(\Omega_{X_i}, \mathbb{R}^d)$ and $(\Omega_{X_j}, \mathbb{R}^d)$ are conjugate.

Since $M(n,L)$ is finite, we can take the same constant $s_0 > 0$ and R_π of Lemmas 23 and 25 respectively, associated to each π_i. Fix $0 < \varepsilon < 1$. Let $R > \sup\{s_0, R_\pi + \varepsilon, 17L\}$ be sufficiently large such that Lemma 6 and Lemma 29 are applicable to R-patches of X, and such that Lemma 23 is applicable to ε and each π_i.

Consider the patch $\mathtt{P} = B_R(0) \cap X$, and $v_1, \ldots, v_N \in X_{\mathtt{P}}$ such that for every $v \in X_{\mathtt{P}}$, there exist $1 \leq i \leq N$ and $u \in \mathbb{R}^d$ satisfying $V_{\mathtt{P},v} \cap X = (V_{\mathtt{P},v_i} \cap X) + u$. Roughly speaking, every set of the kind $V_{\mathtt{P},v} \cap X$ is a translated of some set $V_{\mathtt{P},v_i} \cap X$. Since $R > R_1$, Lemma 6 ensures $N \leq c(L)$.

For every $1 \leq j \leq N$, let $w_{j,1}, \ldots, w_{j,m_j}$ be return vectors of $V_{\mathtt{P},v_j} \cap X$, chosen in order that for every return vector w of $V_{\mathtt{P},v_j} \cap X$, there exists $1 \leq i \leq m_j$ such that $\mathtt{P}_{n,w,v_j}$ is equal to $\mathtt{P}_{n,w_{j,i},v_j} =: \mathtt{P}_{w_{j,i}}$. Since $R > R_1$, Lemma 29 implies that $m_j \leq c(n,L)$, for every $1 \leq j \leq N$. Therefore, the collection

$$\mathcal{F} = \{\mathtt{P}_{w_{j,l}} : 1 \leq l \leq m_j,\ 1 \leq j \leq N\}$$

contains at most $c(L)c(n,L)$ elements.

We define $R' = (L^n - 1)R - \varepsilon - 4LR$. The choice of n ensures that $R' > 0$.

For every $1 \leq i \leq M(n,L)+1$, we define the following relation on $\mathcal{F}$:

$\mathtt{P}_{w_{j,l}} \leftrightarrow_i \mathtt{P}_{w_{k,m}}$ if and only if for every $X', X'' \in \Omega_X$ such that $X' \cap B_{L^n R}(0) = \mathtt{P}_{w_{j,l}}$ and $X'' \cap B_{L^n R}(0) = \mathtt{P}_{w_{k,m}}$, there exist $v \in B_{2\varepsilon}(0)$ and $w \in B_{4LR}(0)$ such that $\pi_i(X') \cap B_{R'}(0) = (\pi_i(X'') + v + w) \cap B_{R'}(0)$.

Since $L^n R - s_0 \geq (L^n - 1)R \geq R$, from Lemma 23 it follows this relation is reflexive, so non empty. Since the cardinal of $\mathcal{F}$ is bounded by $c(L)c(n,L)$, there are at most $M(n,L)$ different relations of this kind. So, there exist $1 \leq i < j < M(n,L)+1$ such that $\leftrightarrow_i = \leftrightarrow_j$.

In the sequel, we will prove that $(\Omega_{X_i}, \mathbb{R}^d)$ and $(\Omega_{X_j}, \mathbb{R}^d)$ are conjugate. For that, it is sufficient to show that if $Y, Z \in \Omega_X$ are such that $\pi_i(Y) = \pi_i(Z)$ then $\pi_j(Y) = \pi_j(Z)$.

Let Y and Z be two Delone sets in Ω_X such that $\pi_i(Y) = \pi_i(Z)$. Without loss of generality, we can suppose that 0 is an occurrence of $\mathtt{P}$ in Y and in $Z - u_0$, where u_0 is some point in $B_{4LR}(0)$. The patches of Y and Z are translated of the patches of X. This implies there exist $1 \leq q_0, r_0 \leq N$ such that

$$Y \cap B_{L^n R}(0) = \mathtt{P}_{w_{q_0,l_0}} \text{ and } (Z - u_0) \cap B_{L^n R}(0) = \mathtt{P}_{w_{r_0,k_0}},$$

for some $1 \leq l_0 \leq m_{q_0}$ and $1 \leq k_0 \leq m_{r_0}$.

It is possible to show that $\mathtt{P}_{w_{q_0,l_0}} \leftrightarrow_i \mathtt{P}_{w_{r_0,k_0}}$ and $\mathtt{P}_{w_{q_0,l_0}} \leftrightarrow_j \mathtt{P}_{w_{r_0,k_0}}$ for R sufficiently large (see Claim 1 in the proof of [CDP, Thm. 12]).

Let s be any other occurrence of $\mathtt{P}$ in Y. Repeating the same argument for $Y + s$ and $Z + s$, we deduce there exist $u_s \in B_{4LR}(0)$ and $1 \leq q_s, r_s \leq N$ such that

$$(Y + s) \cap B_{L^n R}(0) = \mathtt{P}_{w_{q_s,l_s}} \text{ and } (Z + s - u_s) \cap B_{L^n R}(0) = \mathtt{P}_{w_{r_s,k_s}},$$

for some $1 \leq l_s \leq m_{q_s}$ and $1 \leq k_s \leq m_{r_s}$. Then we get $\mathtt{P}_{w_{q_s,l_s}} \leftrightarrow_j \mathtt{P}_{w_{r_s,k_s}}$. This implies there exist $t_s \in B_{2\varepsilon}(0)$ and $w_s \in B_{4LR}(0)$ such that

$$\pi_j(Y + s) \cap B_{R'}(0) = (\pi_j(Z + s - u_s) + t_s + w_s) \cap B_{R'}(0).$$

Showing that $w_s - u_s + t_s$ does not depend on s (see Claim 2 in the proof of [CDP, Thm. 12]), we get there exists $y \in \mathbb{R}^d$ such that for every occurrence s of $\mathtt{P}$ in Y,

$$\begin{aligned} \pi_j(Y + s) \cap B_{R'}(0) &= (\pi_j(Z + s) + y) \cap B_{R'}(0), \text{ and then} \\ \pi_j(Y) \cap B_{R'}(s) &= (\pi_j(Z) + y) \cap B_{R'}(s). \end{aligned}$$

The diameter of the Voronoï cells of $\mathtt{P}$ is less than $4LR$ (see 2.1), which is less than R'. Hence,

$$\pi_j(Y) = \pi_j(Z) + y.$$

We conclude with Propositions 27 and 28. □

5.2. Factors on groups and cocycles

Cocycles and cohomological equations play an important role in the study of factors dynamical systems, time change for flows orbit equivalence, ... We adapt this notion to the context of Delone system $(\Omega, \mathbb{R}^d)$. Let G denotes the group $\mathbb{R}^m$ or $\mathbb{T}^m = \mathbb{R}^m/\mathbb{Z}^m$. A *continuous G-cocycle* is a continuous function $\alpha\colon \Omega \times \mathbb{R}^d \to G$ so that

$$\alpha(Y, v+w) = \alpha(Y, v) + \alpha(Y+v, w) \qquad \text{for all } Y \in \Omega, v, w \in \mathbb{R}^d.$$

An important question which appears in many problems, is to known if the *cohomological equation*

$$\alpha(Y, v) = \psi(Y+x) - \psi(Y)$$

has a measurable, continuous solution $\psi : \Omega \to G$. This solution is called a *transfer function* and if it exists, α is called a *coboundary*.

In Section 5.2.2 we will give a necessary and sufficient condition to find solutions to the cohomological equation for linearly repetitive Delone systems. We will focus on *transversally locally constant* cocycle α, *i.e.*: if there exists $r, R > 0$ such that for any $Y, Y' \in \Omega$ and $x \in B_R(0)$,

$$\text{if } Y \cap B_R(0) = Y' \cap B_R(0) \text{ then } \alpha(Y, x) = \alpha(Y', x).$$

More generally a cocycle α is *transversally Hölder* if there exist constants $K > 0$ and $\delta \in (0, 1)$ such that for all $r > 0$, $Y, Y' \in \Omega$ and $x \in B_r(0)$,

$$\text{if } Y \cap B_R(0) = Y' \cap B_R(0) \text{ then } |\alpha(Y, x) - \alpha(Y', x)| \leq K r^{-\delta}.$$

5.2.1. Examples of cohomological equations. Let us see first some dynamical problems where the cohomological equation appears.

Let us denote by $\langle ., . \rangle$ the usual inner product in $\mathbb{R}^d$ and μ be an ergodic $\mathbb{R}^d$ invariant probability measure on the hull Ω. A vector $\lambda \in \mathbb{R}^d$ is a *measurable eigenvalue* of the system $(\Omega, \mathbb{R}^d)$ if there exists a measurable function $\psi\colon \Omega \to \mathbb{S}^1$ such that

$$\psi(Y+v) = e^{2i\pi\langle \lambda, v\rangle}\psi(Y) \quad \text{for all } v \in \mathbb{R}^d \text{ and } \mu - a.e.\ Y \in \Omega.$$

If the function ψ is continuous, then λ is called a *continuous eigenvalue*. The map $(Y, v) \mapsto e^{2i\pi\langle \lambda, v\rangle}$ is a $\mathbb{S}^1$-cocycle over $(\Omega, \mathbb{R}^d)$. Then passing in additive notation $\mathbb{T}^1$, we have λ is a measurable (resp. continuous) eigenvalue of $(\Omega, \mathbb{R}^d)$ if and only if there is a measurable (resp. continuous) solution $\psi : \Omega \to \mathbb{T}^1$ to the cohomological equation

$$\langle \lambda, v\rangle = \psi(Y+v) - \psi(Y) \bmod \mathbb{Z}.$$

A continuous eigenvalue gives then a topological factor on the closure of an orbit in the one-dimensional torus $\mathbb{T}^1$. More generally, one can consider the closure $\mathbf{O}$ of an orbit of a n-rotations on the n-torus $\mathbb{T}^n$, $n \leq d$, that are factors of the system $(\Omega, \mathbb{R}^d)$. More precisely, take $\theta = (\theta_1, \ldots, \theta_n) \in \mathbb{R}^n$ and let $\mathcal{A} : \mathbb{R}^d \times \mathbb{T}^n \to \mathbb{T}^n$ be the continuous action defined by

$$\mathcal{A}(v, x) = x + [v, \theta] \qquad \text{where } [v, \theta] = (v_1\theta_1, \ldots, v_n\theta_n).$$

The map $(Y, v) \mapsto [v, \theta]$ is a $\mathbb{T}^n$-cocycle over $(\Omega, \mathbb{R}^d)$. It is standard to show that the system $(\mathbf{O}, \mathcal{A})$ is a topological factor of $(\Omega, \mathbb{R}^d)$ if and only if there exists a continuous solution $\psi : \Omega \to \mathbb{T}^n$ to the cohomological equation

$$[v, \theta] = \psi(Y + v) - \psi(Y).$$

5.2.2. Characterization of continuous coboundary. A seminal work for the characterization of continuous eigenvalues of symbolic systems given by a primitive substitution, is in [H]. The authors of [CDHM, BDM1] generalize these results to the linearly recurrent symbolic systems and to finite rank systems in [BDM2]. An extension to $\mathbb{Z}^d$-action on a Cantor set is presented in [CGM]. We present here a part of the results in [C] that treat only continuous cocycles and generalizes the results of [CGM].

For a box decomposition $\mathcal{B} = \{C_i[D_i]\}_{i=1}^t$ (see Section 3.2.1), a *first retrun vector to* $C = \cup_i C_i$ is a vector $v \in \mathbb{R}^d$ with label $(i, j) \in \{1, \dots, t\}^2$, such that

$$C_i - v \cap C_j \neq \emptyset \text{ and } C_i[D_i] \cap C_j[D_j] \neq \emptyset.$$

We denote by $\mathcal{F}$ the set of first return vectors to C associated with $\mathcal{B}$, and by $C(v) = C_i \cap (C_j + v)$ for a return vector v with label (i, j).

A tower system $(\mathcal{B}_n = \{C_{n,i}[D_{n,i}]\}_{i=1}^{t_n})_n$ is *well distributed* if it satisfies the properties i)–iii) in Theorem 13 and moreover for every $n \geq 0$, and every first return vector $v \in \mathcal{F}_n$ with label (i, j) there are x and x' in $D_{n+1,1}$ such that for $X \in \bigcap_n C_n$, $X - x \in C_{n,i}$ and $X - x' \in C_{n,j}$ and $v = x - x'$.

It is straightforward to check that this extra condition holds when each $D_{n+1,i}$ is big enough: more precisely when for any $n \geq 0$

$$r_{\text{int}}(\mathcal{B}_{n+1}) \geq (R_{\text{rec}}(\mathcal{B}_n) + R_{\text{ext}}(\mathcal{B}_n))\, L \geq M_X(R_{\text{rec}}(\mathcal{B}_n) + R_{\text{ext}}(\mathcal{B}_n)). \tag{5.8}$$

For a linearly repetitive Delone set X, it is direct to check that for a constant K big enough, the tower system given by Theorem 13, satisfies inequality (5.8). Thus any linearly repetitive Delone system admits a well-distributed tower system. In the following $|\cdot|$ denotes the usual distance to the origin when $G = \mathbb{R}^m$ or $\mathbb{T}^m$.

Theorem 30 ([C]). *Let X be a linearly repetitive Delone set in $\mathbb{R}^d$, G be the group $\mathbb{R}^m$ or $\mathbb{T}^m$, α be a continuous G-cocycle over $(\Omega, \mathbb{R}^d)$, and $(\mathcal{B}_n)_{n\geq 0}$ be a well-distributed tower system. Then α is a tansversally Hölder coboundary with continuous transfer function if and only if the series*

$$\sum_{n\geq 0} \sup_{\substack{v\in\mathcal{F}_n \\ \omega\in C_n(v)}} |\alpha(\omega, v)|$$

converges, where each $\mathcal{F}_n$ denotes the set of first return vectors associated with $\mathcal{B}_n$.

In [C] appears also similar necessary conditions for a cocycle to be a coboundary on a general Delone system (without the assumption of linear repetitivity).

5.2.3. Characterization of the measurable eigenvalues. To be more complete on the problem of eigenvalues, let us mention that a characterization of measurable eigenvalues of linearly recurrent Cantor system is given in [BDM1] and measurable coboundary for linearly repetitive Delone systems in [C0].

Theorem 31 ([C0]). *Let $(\Omega, \mathbb{R}^d)$ be a linearly repetitive Delone system, μ be the unique invariant measure, G be the group $\mathbb{R}^m$ or $\mathbb{T}^m$, α be a transversally locally constant G-cocycle over $(\Omega, \mathbb{R}^d)$, and $(\mathcal{B}_n)_{n\geq 0}$ be a tower system well-distributed. Then the following are equivalent.*

1. *The series $\sum_{n\geq 0} \sup_{\substack{v\in\mathcal{F}_n \\ \omega\in C_n(v)}} |\alpha(\omega, v)|^2$ converges, where each $\mathcal{F}_n$ denotes the set of first return vectors associated with $\mathcal{B}_n$.*
2. *There exists a measurable function $\psi\colon \Omega \to G$ such that for μ-a-e $X \in \Omega$,*
$$\alpha(X, v) = \psi(X - v) - \psi(X), \qquad \text{for all } v \in \mathbb{R}^d.$$
Moreover $\psi \in L^2(\Omega, \mathbb{R}^m, \mu)$ when $G = \mathbb{R}^m$.

6. Bi-Lipschitz equivalence to a lattice

Let X_1 and X_2 be two Delone sets in $\mathbb{R}^d$. We say that they are *bi-Lipschitz equivalent* if there exists a homeomorphism $\phi : X_1 \to X_2$ and a constant $\Delta \geq 1$ such that $\forall x, x' \in X, x \neq x'$
$$\frac{1}{\Delta} \leq \frac{\|\phi(x) - \phi(x')\|}{\|x - x'\|} \leq \Delta.$$
The map ϕ is then called a *bi-Lipschitz homeomorphism* between X_1 and X_2.

The problem to know whether two Delone sets are bi-Lipschitz equivalent was raised by Gromov in [Gro93], and boiled down in Toledo's review [Tol] to the following question for the two-dimensional Euclidean space: *Is every Delone set in $\mathbb{R}^2$ bi-Lipschitz equivalent to $\mathbb{Z}^2$?* Counterexamples to this question were given independently by Burago and Kleiner [BK] and McMullen [McM]. Moreover, McMullen also showed that when relaxing the bi-Lipschitz condition to a Hölder one, all Delone set (with or without finite local complexity) in $\mathbb{R}^d$ are equivalent. Later, Burago and Kleiner [BK1] gave a sufficient condition for a Delone set to be bi-Lipschitz equivalent to $\mathbb{Z}^2$ and asked the following question: *If one forms a Delone set in the plane by placing a point in the center of each tile of a Penrose tiling, is the resulting set bi-Lipschitz equivalent to $\mathbb{Z}^2$?* They studied the more general question of knowing whether a Delone set arising from a cut-and-project tiling is bi-Lipschitz equivalent to $\mathbb{Z}^2$ (recall that the Penrose tiling is also a cut-and-project tiling [Bru]) and solved it in some cases that do not include the case of Penrose tilings, thus leaving the former question open. Recently, Solomon [Solo] gave a positive answer for Penrose tiling by using the fact that it can be constructed using substitutions. In fact, Solomon proved that each Delone set arising from a primitive self-similar tiling in $\mathbb{R}^2$ is bi-Lipschitz to $\mathbb{Z}^2$.
The following result was proved in [ACG1].

Theorem 32. *Every linearly repetitive Delone set in $\mathbb{R}^d$ is bi-Lipschitz equivalent to $\mathbb{Z}^d$.*

Notice that Theorem 32 is trivial when the dimension $d = 1$ since, in this case, every Delone set (with no extra assumptions) is bi-Lipschitz equivalent to $\mathbb{Z}$. As an application of the work of Laczkovich [L], Solomon in [Solo] showed also that for every self-similar tiling of $\mathbb{R}^d$ of Pisot type there is a *bounded displacement* between its associated Delone set X and $\beta\mathbb{Z}^d$ for a $\beta > 0$ (i.e., there is a bijection $\phi\colon X \to \beta\mathbb{Z}^d$ such that $\Phi - Id$ is bounded).

The strategy of the proof of Theorem 32 is the following. First consider the easy case where all the Voronoï cells V of a Delone set X have a unit volume. Thus any finite union of N Voronoï cells meet at least N unit squares, and conversely N unit squares meet at least N Voronoï cells. So by the transfinite form of Hall's marriage lemma, there exists a bijection between the collection of Vornoï cells and the units squares, so that any cell intersects its image. This defines a map $\phi\colon X \to \mathbb{Z}^d$ such that $\phi - Id$ is bounded.
For the general case, we need to consider the measurable function $f\colon \mathbb{R}^d \to \mathbb{R}$ defined by

$$f(x) = \sum_{y:x\in V_y} \frac{1}{\operatorname{vol} V_y} \qquad x \in \mathbb{R}^d,$$

where V_y denotes the Voronoï cell of the point $y \in X$. If $\phi : \mathbb{R}^d \to \mathbb{R}^d$ is a bi-Lipschitz map so that its Jacobian determinant is f, standard calculus show us that the image $\phi(V)$ of any Voronoï cell V of X has volume 1. The proof of Theorem 32 consists then to generalize to all dimension d a sufficient condition given by Burago and Kleiner [BK1] in dimension 2 to solve the equation $\det \mathrm{D}\phi = f$ with ϕ an unknown bi-Lipschitz map. This condition involves analytical tools and the density deviation of the points of X with respect to its average. This last point is controlled by the Lagarias and Pleasants Theorem 16.

References

[AC] J. Aliste Prieto, D. Coronel, *Tower systems for linearly repetitive Delone sets*, Ergod. Th. and Dynam. Sys. (2011) 31 6 pp. 1595 –1618

[ACG] J. Aliste Prieto, D. Coronel, J.-M. Gambaudo, *Rapid convergence to frequency for substitution tilings of the plane.* Comm. Math. Phys. 306 (2011), no. 2, 365–380

[ACG1] J. Aliste Prieto, D. Coronel, J.-M. Gambaudo, *Linearly repetitive Delone sets are rectifiable.* arXiv:1103.5423.

[BBG] J. Bellissard, R. Benedetti, J.-M. Gambaudo, *Spaces of tilings, finite telescopic approximations and gap-labeling.* Comm. Math. Phys. 261 (2006), no. 1, 1–41.

[BG] R. Benedetti, J.-M. Gambaudo, *On the dynamics of $\mathbb{G}$-solenoids. Applications to Delone sets.* Ergodic Theory Dynam. Systems 23 (2003), no. 3, 673–691

[Bes] A. Besbes, *Contributions a l'étude de quelques systèmes quasi-crystallographics*, Ph.D. thesis, Université Pierre et Marie Curie, Paris, 2008.

[BBL] A. Besbes, M. Boshernitzan, D. Lenz, *Delone sets with finite local complexity: linear repetitivity versus positivity of weights.* Preprint

[Bo] M. Boshernitzan, *A condition for unique ergodicity of minimal symbolic flows.* Ergodic Theory Dynam. Systems 12 (1992), no. 3, 425–428.

[BDM1] X. Bressaud, F. Durand, and A. Maass, *Necessary and sufficient conditions to be an eigenvalue for linearly recurrent dynamical Cantor systems*, J. London Math. Soc. (2) 72 (2005), no. 3, 799–816.

[BDM2] X. Bressaud, F. Durand, and A.Maass, *On the eigenvalues of finite rank Bratteli–Vershik dynamical systems*. Ergodic Theory Dynam. Systems 30 (2010), no. 3, 639–664.

[Bru] N.G. de Bruijn, *Algebraic theory of Penrose's nonperiodic tilings of the plane. I, II.* Nederl. Akad. Wetensch. Indag. Math. 43 (1981), no. 1, 39–52, 53–66.

[BK] D. Burago, B. Kleiner, *Separated nets in Euclidean space and Jacobians of bi-Lipschitz maps.* Geom. Funct. Anal. 8 (1998), no. 2, 273–282.

[BK1] D. Burago, B. Kleiner, *Rectifying separated nets.* Geom. Funct. Anal. 12 (2002), no. 1, 80–92.

[CD] M.I. Cortez, F. Durand, *Self-similar tiling systems, topological factors and stretching factors*, Discrete Comput. Geom. 40 (2008), no. 4, 622–640.

[CDHM] M.I. Cortez, F. Durand, B. Host, A. Maass, *Continuous and measurable eigenfunctions of linearly recurrent dynamical Cantor systems.* J. London Math. Soc. (2) 67 (2003), no. 3, 790–804.

[CDP] M.I. Cortez, F. Durand, S. Petite, *Linearly repetitive Delone systems have a finite number of nonperiodic Delone system factors.* Proc. Amer. Math. Soc. 138 (2010), no. 3, 1033–1046.

[CGM] M.I. Cortez, J.M. Gambaudo, and A. Maass, *Rotation topological factors of minimal $\mathbb{Z}^d$-actions of the Cantor set*, Trans. Amer. Math. Soc. 359 (2007), no. 5, 2305–2315.

[C] D. Coronel, *The cohomological equation over dynamical systems arising from Delone sets.* Ergod. Th. and Dynam. Sys. 31 (2011), no. 3, 807–833.

[C0] D. Coronel, *Ecuaciones cohomológicas sobre espacios de embaldosados*, Ph.D. thesis, Univ. de Chile, Univ. de Nice-Sophia Antipolis, Santiago, Nice, 2009.

[DL] D. Damanik, D. Lenz, *Linear repetitivity I., Uniform subadditive ergodic theorems and applications*, Discrete and Computational Geometry, 26 (2001), 411–428.

[DK] F.M. Dekking, M. Keane, *Mixing properties of substitutions*, Zeit. Wahr. **42** (1978), 23–33

[Du] F. Durand, *Contributions à l'étude des suites et systèmes dynamiques substitutifs*, Ph.D. thesis, Univ. de la Méditerrannée, Aix-Marseille II, Marseille, 1996.

[Du0] F. Durand, *A characterization of substitutive sequences using return words*, Discrete Mathematics 179 (1998), 89–101.

[Du1] F. Durand, *Linearly recurrent subshifts have a finite number of non-periodic subshift factors*, Ergod. Th. and Dynam. Sys. 20 (2000), 1061–1078.

[Du2] F. Durand, *Corrigendum and addendum to: Linearly recurrent subshifts have a finite number of non-periodic subshift factors*, Ergod. Th. and Dynam. Sys. 23 (2003), 663–669.

[DHS] F. Durand, B. Host, C. Skau, *Substitutional dynamical systems, Bratteli diagrams and dimension groups*, Ergod. Th. and Dynam. Sys. 19 (1999), 953–993.

[Gh] Étienne Ghys, *Laminations par surfaces de Riemann*, Dynamique et géométrie complexes (Lyon, 1997), vol. 8, Soc. Math. France, Paris, 1999, pp. 49–95.

[Gro93] M. Gromov, *Asymptotic invariants of infinite groups.* Geometric group theory, Vol. 2 (Sussex, 1991), 1–295, London Math. Soc. Lecture Note Ser., 182, Cambridge Univ. Press, Cambridge, 1993.

[H] Host, B. *Valeurs propres des systèmes dynamiques définis par des substitutions de longueur variable.* Ergodic Theory Dynam. Systems 6 (1986), no. 4, 529–540.

[HRS] C. Holton, C. Radin, L. Sadun, *Conjugacies for tiling dynamical systems*, Comm. Math. Phys. 254 (2005), 343–359.

[L] M. Laczkovich, *Uniformly spread discrete sets in* $\mathbb{R}^d d$, Journal of the London Mathematical Society. Second Series 46 (1992), 39–57.

[La] J.C. Lagarias, *Geometric models for quasicrystals I. Delone sets of finite type.* Discrete Comput. Geom. 21 (1999), no. 2, 161–191.

[LP1] J.C. Lagarias, P.A.B. Pleasants, *Local complexity of Delone sets and crystallinity.* Dedicated to Robert V. Moody. Canad. Math. Bull. 45 (2002), no. 4, 634–652.

[LP2] J.C. Lagarias, P.A.B. Pleasants, *Repetitive Delone sets and quasicrystals*, Ergod. Th. and Dynam. Sys. 23 (2003), 831–867.

[Le02] D. Lenz, *Uniform ergodic theorems on subshifts over a finite alphabet.* Ergodic Theory Dynam. Systems 22 (2002), no. 1, 245–255.

[Le] D. Lenz, *Aperiodic linearly repetitive Delone sets are densely repetitive.* Discrete Comput. Geom. 31 (2004), no. 2, 323–326

[McM] C.T. McMullen, *Lipschitz maps and nets in Euclidean space.* Geom. Funct. Anal. 8 (1998), 304–314.

[MH] M. Morse, G. Hedlund, *Symbolic dynamics II. Sturmian trajectories.* Amer. J. Math. 62, (1940). 1–42.

[Pe] K. Petersen, *Factor maps between tiling dynamical systems*, Forum Math. 11 (1999), 503–512.

[RS] C. Radin, L. Sadun, *Isomorphism of hierarchical structures*, Ergod. Th. and Dynam. Sys. 21 (2001), 1239–1248.

[Solo] Y. Solomon, *Substitution Tilings and Separated Nets with Similarities to the Integer Lattice*, Israel J. Math. 181 (2011), no. 1, 445–460.

[So1] B. Solomyak, *Dynamics of self-similar tilings*, Ergod. Th. and Dynam. Sys. 17 (1997), 695–738.

[So2] B. Solomyak, *Nonperiodicity implies unique composition for self-similar translationally finite tilings*, Discrete Comput. Geom. 20 (1998), 265–279.

[So3] B. Solomyak, *Spectrum of dynamical systems arising from Delone sets*, In: Quasicrystals and discrete geometry (Toronto, ON, 1995), Providence, RI: Am. Math. Soc., 1998, 265–275.

[Tol] D. Toledo, *Geometric group theory, 2: Asymptotic invariants of finite groups by M. Gromov*. Bull. Amer. Math. Soc. 33 (1996), 395–398.

José Aliste-Prieto and Daniel Coronel
Departamento de Matemáticas
Universidad Andres Bello
Republica 220, 2do. piso
Santiago, Chile

María Isabel Cortez
Departamento de Matemática
y CC. de la Universidad de
Santiago de Chile,
Av. Libertador Bernardo O'Higgins 3363
Santiago, Chile
e-mail: mcortez@usach.cl

Fabien Durand and Samuel Petite
Laboratoire Amiénois de Mathématique
Fondamentale et Appliquée, CNRS-UMR 7352
Université de Picardie Jules Verne
33 rue Saint Leu
F-80039 Amiens Cedex, France
e-mail: fabien.durand@u-picardie.fr
samuel.petite@u-picardie.fr

Progress in Mathematics, Vol. 309, 223–257
© 2015 Springer Basel

Tilings with Infinite Local Complexity

Natalie Priebe Frank

Abstract. This is a chapter surveying the current state of our understanding of tilings with infinite local complexity. Such tilings can arise when tiles have infinitely many possible adjacencies, infinitely many shapes, or infinitely many labels. Our main requirement is that the set of tiles used to construct tilings should be compact.

We consider tilings constructed in a number of ways, including the hierarchical methods of self-similarity, substitution, and fusion. We show how to adapt the standard toolbox for tilings with finite local complexity and suggest definitions for the concepts of fault lines and complexity functions. Three examples with infinite local complexity of distinctly different origin are fully analyzed using the tools and techniques contained in this chapter. We conclude with some important classes of open questions about tiling spaces with infinite local complexity.

Mathematics Subject Classification (2010). Primary 52C23; Secondary 37B50.

Keywords. Self-similarity, fusion tilings, tiling dynamical systems.

1. Introduction

Most of the literature on tiling spaces and their dynamical systems has focused on those with finite local complexity (FLC). In this paradigm there is a finite set $\mathcal{P}$ of tiles called 'prototiles', congruent copies of which are used to cover the plane (or $\mathbb{R}^d$) without gaps or overlaps. Moreover, the adjacencies between tiles are restricted so that there are only finitely many two-tile configurations. If there can be infinitely many two-tile configurations in a tiling, then that tiling is said to have infinite local complexity (ILC).

When tilings are looked at from a physical perspective it makes sense to consider not just individual tilings but rather spaces whose elements are tilings that share some common properties. These tiling spaces are given a metric topology where the distance between two tilings is defined by how similar they are in balls

Some work presented here was done in collaboration with Lorenzo Sadun and Ian Putnam, both of whom the author thanks for their hospitality and many illuminating discussions.

around the origin (see Section 2.2 for precision). When there are only finitely many two-tile configurations that are found in any tiling of a tiling space, we say that the tiling space itself has finite local complexity; otherwise, it has infinite local complexity. FLC tiling spaces have been the standard objects used to model the atomic structure of crystals and quasicrystals and have proved quite effective in the study of statistical properties, diffraction patterns, and energy spectra of aperiodic solids.

A tiling space that is of finite local complexity can be homeomorphic to one with infinite local complexity [20]. Thus finite local complexity is not a topological invariant and should not be considered an intrinsic property when topological methods are used to study aperiodic tilings.

Examples of tilings with infinite local complexity have appeared sporadically [6, 9, 10, 16, 18, 21], and it is increasingly clear that the class is not as unnatural as previously imagined. Moreover, most of the 'usual' FLC tools and techniques can be used in the ILC case, and one of the goals of this chapter is to explain exactly how to adapt the existing machinery. We take as fundamental the requirement that prototiles come from *compact*, not necessarily finite, sets. This means that both the 'supports' of the tiles (i.e., their underlying sets in $\mathbb{R}^d$) and the 'labels' of the tiles (which are used to distinguish tiles with congruent supports) must come from compact sets. We will see that this fundamental requirement means that ILC tiling spaces are compact (see Section 2.3). We delay formal definitions until Section 2 and provide some informal examples now.

1.1. Introductory examples

Since one-dimensional tiles are closed intervals, any tiling made from a finite number of interval lengths with a finite number of labels must have finite local complexity. So in order to have infinite local complexity in one dimension there must be either an infinite label set or an infinite number of lengths (or both).

Example 1. A first example is to allow tiles to take lengths from some closed interval, for instance we could require that $1 \leq \text{length} \leq 3$. We can let the support of a prototile p_x be the interval $[0, x]$, and we can label the tile by its length, x. It is convenient to omit the label when it is possible to tell tiles apart by their supports, as is the case here, but we are including them for consistency with the definitions provided in Section 2. The prototile p_x is formally the pair $([0, x], x)$, and the prototile set in this example is thus $\mathcal{P} = \{p_x, x \in [1, 3]\}$. Notice that the set of supports of prototiles is compact in the Hausdorff metric and the set of labels is compact in the usual distance metric in $\mathbb{R}$. Importantly, if a sequence of prototiles has a convergent label sequence, then their supports converge as well. This makes it possible to say that the sequence of prototiles themselves converge.

A tile is simply a translate of p_x by some element $y \in \mathbb{R}$; we write $t = p_x + y = ([y, y + x], x)$. (Note that translation changes the support of a tile but not its label.) We could make a tiling from such tiles in any number of ways, for instance by generating a sequence of random numbers in $[1, 3]$ and laying down tiles of those lengths in any order. With probability one such a tiling will have

infinite local complexity because it has infinitely many different tiles and thus has infinitely many different two-tile patterns.

This example provides a nice test case for computation since it is really different than the standard FLC situation, but still quite simple. Example 5 consists of a hierarchical tiling space based on a prototile set derived from $\mathcal{P}$. We introduce its construction in Section 4.1 and give it a thorough analysis in Section 6.1.

Example 2. For another one-dimensional example we take a single interval length for the supports but allow for infinitely many labels. Suppose that the support of every prototile is $[0, 1]$, but that each prototile takes a label from some compact label set $\mathcal{L}$. For concreteness, let $\mathcal{L} = S^1 = \mathbb{R}/\mathbb{Z}$, the unit circle. We cannot tell two prototiles apart by their supports, so the label tells us when two of them are different, and the distance between their labels tells us how different they are.

An interesting way to construct a tiling from this prototile set is to fix an element $\alpha \in S^1$ consider the sequence of labels $x + n\alpha \mod 1$ for any $x \in \mathbb{R}$. If α is irrational then the tilings generated by this label sequence will have infinite local complexity since $\{n\alpha \mod 1\}$ is infinite (and in fact uniformly distributed).

Example 3. A two-dimensional example of an ILC tiling can be constructed from unit squares. Tile the plane in rows of tiles, but offset each row from the next by a randomly chosen number in [0,1]. With probability one, the result will be an ILC tiling. A non-random variation on this theme is to base the offsets on some fixed irrational number α. Lay the first row of squares along the x-axis with a vertex at the origin. Place an endpoint of the row at height $y = 1$ at $x = \alpha$, and the endpoint of the row at height $y = n$ at $x = n\alpha$. Since $\{n\alpha \mod 1\}$ is uniformly distributed in $[0, 1]$, the offsets between rows will be too and in this way form an ILC tiling of the plane.

Example 4. A well-known example that has infinite local complexity up to translations is the pinwheel tiling. Pinwheel tiles appear in infinitely many orientations in any individual pinwheel tiling and so there are not finitely many different two-tile patches that are translates of one another. This is a borderline case, however: the tiles fit together in finitely many ways even though these allowed configurations appear in infinitely many orientations. It is sometimes useful, then, to consider the pinwheel tiling space to be of finite local complexity by allowing rotations along with translations.

1.2. Ways infinite local complexity arises

In higher dimensions there are many natural examples of tilings with infinite local complexity. For instance, the atomic structure of an ideal crystal is modeled by a lattice of points, but the atoms in an actual crystal appear within a certain tolerance of that lattice. A standard perspective to let the atomic structure generate a tiling, either by using the atoms as tile vertices or, by taking the Voronoï tessellation of the set of atomic locations, or by some other method. In the case of an ideal crystal, all methods yield periodic tilings with patches of tiles forming

unit cells. In the case of an actual crystal, however, these tiles will be deformed within a certain tolerance and we will have an infinite number of tile shapes.

Infinite local complexity has long been known to arise even when there are a finite number of tile shapes. There are tilings with a finite number of tile types inside of which 'fault lines' develop. Defined formally in Section 5.1, a fault line separates a tiling into half-tilings that can slide parallel to the fault line to produce new tilings from the tiling space. The presence of fault lines often result in an infinite number of local adjacencies. If we were to encode adjacency information as labels for the tiles, then we would have infinitely many labels. If a tiling of $\mathbb{R}^2$ has only a finite number of tile shapes up to Euclidean motions, it is proved in [16] that there are only two ways ILC can appear: either along a fault line or along a fault circle. The former case requires tiles that have a straight edge somewhere, while the latter requires tiles with an edge that is a circular arc of some given radius.

1.3. Outline of this chapter

Section 2 contains the details on how we conceive of tiles, tilings, tiling spaces, and the tiling metric in the presence of infinite local complexity. Our definitions coincide with those for finite local complexity tilings when that condition is satisfied.

Section 3 addresses basic analysis of ILC tiling dynamical systems. The translation dynamical system is defined and we explain what minimality, repetitivity, and expansivity mean in this context. The notion of 'cylinder sets' is adapted from symbolic and FLC dynamics, and we show how to deal with some subtle yet important details that impact how they are used. We show how to think about translation-invariant measures and their relationship to patch frequency. Finally we discuss how to generalize the notions of entropy and complexity to this situation.

Tilings with a hierarchical structure generated by substitution or fusion are the topic of Section 4. The construction methods adapt pretty much directly from the FLC case, except care must be taken to preserve compactness of supertile sets. Transition matrices, so useful in frequency computations for FLC self-similar and fusion tilings, need to be dealt with as transition maps instead. The idea of recognizability takes little work to adapt to the ILC case, but primitivity requires some care.

Existing results on ILC tiling spaces are collected into Section 5. We give a 'fault lines' a proper definition, and since they are not topologically invariant we introduce the related idea of 'fractured' tiling spaces. The effect of fault lines and fractures on the topological spectrum is explained before we move on to results specific to the hierarchical tilings case. The fact that primitivity continues to imply minimality is proved and conditions are given that make the converse true as well. We also explain how to think about the invariant measures for fusion systems. In the special case of fusion tilings with strictly finite supertile sets we show the similarity to FLC fusion tilings. Finally we tell everything that is currently known

on the important question "When is an ILC tiling space homeomorphic to an FLC tiling space?"

In Section 6 we apply our toolbox to three different examples. A point of interest that does not appear elsewhere in the literature is how to see certain tilings ('direct product variations') as projections of stepped, branched surfaces in higher dimensions and how that can give rise to infinite local complexity.

The paper concludes with two main categories of questions about tilings with infinite local complexity. One of these has already been mentioned, the question of when an ILC tiling space is homeomorphic, or even topologically conjugate to, an FLC tiling space. The other type is about how the geometric and combinatorial aspects of the tiles or tilings affect the dynamical, measure-theoretic, or topological properties of their tiling spaces.

2. Compact tiling spaces

2.1. Tiles, patches, and tilings

There are two main ingredients for tiles in a tiling with infinite local complexity: supports and labels. The support is the underlying set in $\mathbb{R}^d$ and the label can be thought of as distinguishing between tiles that have congruent supports, perhaps by color or by orientation. Often it is convenient to more or less ignore the labels but since they are quite handy we include them as a fundamental part of our definition. The support and label sets must work together in a precise way in order to define a coherent prototile set that can be used to construct infinite tilings via translation.

Let $\mathcal{S}$ denote a set of subsets of $\mathbb{R}^d$, each of which is a topological disk containing the origin in its interior. Assume $\mathcal{S}$ is a compact metric space under the Hausdorff metric, in which case $\mathcal{S}$ can serve as a set of prototile supports. Let $\mathcal{L}$ be another compact metric space, to be used as the prototile label set. Let $sp : \mathcal{L} \to \mathcal{S}$ be a continuous surjection called the *support map* that assigns to each label a set in $\mathbb{R}^d$ that serves as the physical tile itself.

Definition 2.1. A *prototile* is a pair $p = (S, l)$, where $S \in \mathcal{S}$, $l \in \mathcal{L}$, and $S = sp(l)$. We call S the *support* and l the *label* of p. A prototile set $\mathcal{P}$ is the set of all prototiles associated to a given label set, support set, and support map.

Since the support map sp is continuous we have the property that if a sequence of labels converges in $\mathcal{L}$, their corresponding supports converge in $\mathcal{S}$. This will give us a way to talk about convergence of prototiles and compactness of the prototile set.

The primary action on tiles will be by *translation* by $x \in \mathbb{R}^d$. If $p = (S, l) \in \mathcal{P}$ we define the $\mathcal{P}$-*tile* or just *tile* $t = p - x$ to be the pair $(S - x, l)$. That is, we translate the support of p to a different location but keep the label the same. As for prototiles, tiles have supports and labels; the support of the above tile t is the

set $\text{supp}(t) = S - x$ and the label of t is l. Given an arbitrary tile t, we have support and label maps such that $\text{supp}(t) \subset \mathbb{R}^d$ and $\text{label}(t) \in \mathcal{L}$.

A handy concept in tiling theory is that of the *control point* of a tile $t = p - x$, where $p \in \mathcal{P}$ and $x \in \mathbb{R}^d$, which is defined simply to be the point x. This point represents the location in t of the origin in p and gives us a point of reference for each tile.

Definition 2.2. A finite union of $\mathcal{P}$-tiles whose supports cover a connected region and intersect only on their boundaries is called a *patch*.

We can write $P = \bigcup_{k=1}^{n} t_k$, where $\bigcup_{k=1}^{n} \text{supp}(t_k)$ is connected and $\text{supp}(t_i) \cap \text{supp}(t_j)$ is either empty or contains only boundary points whenever $i \neq j$. Like tiles, patches can be translated and we define $P - x = \bigcup_{k=1}^{n} (t_k - x)$. Two patches are said to be *equivalent* if they are translates of one another.

Definition 2.3. An infinite union of $\mathcal{P}$-tiles whose supports cover the entirety of $\mathbb{R}^d$ and whose pairwise intersections contain only boundary points is called a *tiling* $\mathbf{T}$.

Like patches and tiles, a tiling can be translated by an element $x \in \mathbb{R}^d$ by translating each tile of $\mathbf{T}$ by x. This produces a new tiling we denote by $\mathbf{T} - x$. Precisely, if $\mathbf{T} = \bigcup_{i \in \mathbb{Z}} t_i$ is a tiling expressed as a union of tiles, then we write $\mathbf{T} - x = \bigcup_{i \in \mathbb{Z}} (t_i - x)$, where $t_i - x = (\text{supp}(t_i) - x, \text{label}(t_i))$. This results in an exact copy of the tiling $\mathbf{T}$, except moved so that what was at the point x is now at the origin.

2.2. Tile, patch, and tiling metrics

In order to understand tiling spaces we need to know how to measure the distance between tiles, patches, and tilings. To simplify notation (but not add confusion, we hope) we will use $d(x, y)$ to denote distance where x and y are tiles, patches, or tilings. Each builds on the last.

The distance between two tiles t_1 and t_2 is the maximum of the Hausdorff distance between the supports of the tiles and the difference between the labels:

$$d(t_1, t_2) = \max(d_H(\text{supp}(t_1), \text{supp}(t_2)), d_L(\text{label}(t_1), \text{label}(t_2))). \tag{1}$$

The distance between two patches P_1 and P_2 can be computed provided the tiles are in one-to-one correspondence. Suppose G is the set of all bijections f assigning a tile from P_1 to a tile from P_2. In this case we define

$$d(P_1, P_2) = \min_{f \in G} \{\max_{t \in P_1} \{d(t, f(t))\}\}. \tag{2}$$

Intuitively, we take the bijection that makes the best fit between the two patches and then consider the maximum distance between tiles paired by the bijection. In

the FLC case patches are always matched up by a congruence, usually a translation, in which case the distance is the length of the translation vector. In the ILC case it is necessary to let the tiles move independently from one patch to the other.

The metric for tilings is based on the patch metric and says that two tilings are close if they very nearly agree on a big ball around the origin. For two tilings $\mathbf{T}_1$ and $\mathbf{T}_2$ we define

$$d(\mathbf{T}_1, \mathbf{T}_2) = \inf_{\epsilon>0} \{ \exists P_1 \subset \mathbf{T}_1 \text{ and } P_2 \subset \mathbf{T}_2 \mid B_{1/\epsilon}(0) \subset \operatorname{supp}(P_i) \text{ and } d(P_1, P_2) < \epsilon \} \tag{3}$$

provided such an ϵ exists and is not greater than 1. If there is no such ϵ, or if the infimum is greater than 1, we define the distance between the tilings to be 1.

2.3. Tiling spaces

Rather than trying to study an individual tiling it often makes sense to study all tilings that have certain properties in common. The standard way to do this, motivated by physical applications, is to construct a topological space of tilings.

Definition 2.4. A *tiling space* Ω is a set of tilings of $\mathbb{R}^d$ that is invariant under the action of translation and closed under the topology given by the tiling metric d.

One common way to make a tiling space is by taking the closure of the translational orbit of some fixed tiling $\mathbf{T}$, in which case we write $\Omega_{\mathbf{T}}$. This tiling space is called the *hull* of $\mathbf{T}$.

Theorem 2.5. *Tiling spaces are compact in the metric topology.*

Proof. We establish sequential compactness for patch sets and then extend to tilings. The key to seeing this is to show that the set of all patches contained in a bounded region and having a fixed number n of tiles is compact for every n. Such a set of patches is parameterized by a bounded subset of $\mathcal{P}^n \times \mathbb{R}^{dn}$, where the elements of $\mathbb{R}^{dn}$ are the locations of the control points and thus lie in a bounded region. The individual tiles in any sequence of patches will have convergent subsequences since $\mathcal{P}$ is compact and the tiles lie in a bounded region. We can diagonalize to get a sequence of patches for which all of the individual tiles converge; since each patch in the sequence is connected and the tiles have nonoverlapping boundaries, the limit will have this property as well. Thus every sequence of n-tile patches in a bounded region has a convergent subsequence. Sequential compactness for Ω now follows by finding subsequences of tilings that have convergent sequences of patches covering larger and larger regions around the origin. □

2.4. The transversal $\Xi(\Omega)$ of a tiling space

In definition 2.1 we defined the prototile set as being a representative set of tiles located so that the origin lies in their support at a control point.

Definition 2.6. The *transversal* $\Xi(\Omega)$ of a tiling space Ω is the set of all tilings in Ω with a control point at the origin. Put another way, $\Xi(\Omega)$ is the set of all tilings in Ω containing a prototile.

Every tiling in Ω is the translation of lots of tilings from the transversal. Moreover, every point in the tiling space has a neighborhood that is homeomorphic to an open set in $\mathbb{R}^d$ crossed with an open subset of the transversal.

Much of the work done on FLC tiling spaces uses the transversal in an essential way. For instance, the C^*-algebra of a tiling space is strongly Morita equivalent to the C^*-algebra of its transversal. This means the K-theory of the tiling space can be computed from the transversal. By the gap-labelling theorem, we then understand the possible energy levels that the tiling space can support when considered as an atomic model. The transversal also makes possible the definition of a Laplace–Beltrami operator that holds information on key mechanical properties of solids. This has been studied in the FLC case in, for example, [15]; there is hope that this analysis can be extended to at least some tilings with infinite local complexity.

Thus it is important to understand the structure of the transversal. When a tiling space has finite local complexity, the transversal is always totally disconnected and, under the condition of repetitivity, is a Cantor set. Tilings with infinite local complexity can also have transversals that are Cantor sets, but they can also have more complicated transversals. Lemma 3.2 of [12] states that having a totally disconnected transversal is a topological invariant of tiling spaces.

Lemma 2.7 ([12]). *If two tiling spaces are homeomorphic and one has a totally disconnected transversal, then so does the other.*

The transversal of the pinwheel tiling looks like two Cantor sets, each crossed with a circle. The way to see this is to first imagine a pinwheel tile with the control point at the origin. The set of all tilings that contain this tile will be a Cantor set since distinct tilings are always separated by some amount determined by the closest place on which they differ, yet each tiling is the limit of a sequence of other tilings. Now this Cantor set must be rotated in all amounts to get half the tiling space. The other half of the space is obtained by doing the same thing with the flip of the pinwheel tile we started with. We describe the nature of the transversal for several examples in Section 6.

3. Ergodic theory applied to ILC tiling systems

Since tilings can be used to model the atomic structure of quasicrystals, the statistical, large-scale approach of ergodic theory makes sense: anything happening on a set of measure zero is not physically observable and so can be ignored. *Miles of Tiles* [19] is an exposition of the method that explains the physical motivation for non-physicists. We begin by interpreting fundamental dynamics concepts to our situation.

3.1. Tiling dynamical systems, minimality, repetitivity, and expansivity

Translation provides a natural action of $\mathbb{R}^d$ on Ω that is continuous in the tiling metric and allows us to take a dynamical approach.

Definition 3.1. A *tiling dynamical system* $(\Omega, \mathbb{R}^d)$ is a tiling space Ω along with the action of $\mathbb{R}^d$ by translation.

A dynamical system is said to be *minimal* if the orbit of every tiling under translation is dense. A minimal FLC tiling system has the property that all possible patches of any size can be found in any given tiling $\mathbf{T}$. Since there are many more patches in an ILC system, minimality guarantees that every patch found in any tiling can be arbitrarily well approximated by one from any given tiling $\mathbf{T}$. It is fairly easy to construct a minimal ILC tiling space by using traditional techniques, for instance with substitution as in Section 4.

A tiling $\mathbf{T}$ is said to be *repetitive* if for every patch P that appears in $\mathbf{T}$ and every $\epsilon > 0$, there is an R for which every ball of radius R in $\mathbf{T}$ contains a patch that is within ϵ of P. The orbit closure of $\mathbf{T}$ is a minimal tiling system if $\mathbf{T}$ is repetitive.

A tiling dynamical system is said to be *expansive* if there is a $\delta > 0$ such that whenever $d(\mathbf{T} - x, \mathbf{T}' - x) < \delta$ for all $x \in \mathbb{R}^d$, it means that $\mathbf{T} = \mathbf{T}' - y$ for some $y \in \mathbb{R}^d$ with $|y| < \delta$. In an expansive system, then, the only way for the entire orbits of two tilings to be close is if they were small translates of one another to begin with. FLC tiling spaces, like their cousins the shift spaces, always have expansive dynamical systems. However, infinite local complexity brings us examples of tiling systems that do not have expansive dynamics. Such an example appears as our example 7.

3.2. The Borel topology and cylinder sets

In classical symbolic dynamics it is commonplace to consider the set of all sequences that have a specific symbol or word in a given location, and this set is called a cylinder set. This notion generalizes nicely to the FLC tiling situation, where we need to specify both a patch P and an open set U, such that the cylinder set $\Omega_{P,U}$ is the set of all tilings that contain the patch P in a location designated by U. Two properties of cylinder sets are essential to bring into the ILC situation. First, they generate the metric topology. Second, they can be used to compute the frequency with which the patch P appears throughout the tiling space.

When we have infinitely many different two-tile patches, the cylinder sets based on single patches do not generate the topology. Moreover, it is possible that every individual patch has frequency 0. This means we need to make cylinder sets based on *sets* of patches, for instance, the set of all patches that are within ϵ of some particular patch. To measure frequency accurately we need to define sets of patches that don't contain any 'repeats' up to translation:

Definition 3.2. A set of patches I is said to be *trim* if, for some fixed open set $U \subset \mathbb{R}^d$ and every $\mathbf{T} \in \Omega$, there is at most one patch $P \in I$ and point $x \in U$ for which $P - x \subset \mathbf{T}$.

Thus a trim set does not contain patches that are arbitrarily small translates of one another, or patches that sit in arbitrarily small translates of other patches.

Definition 3.3. Let $U \subset \mathbb{R}^d$ and let I be a set of patches. The *cylinder set* $\Omega_{I,U}$ is the set of all tilings in Ω for which there is some patch $P \in I$ and point $x \in U$ for which $P - x \in \mathbf{T}$.

If I is a trim set with a small enough U, we know that a tiling can only be in the cylinder set via one specific patch P and point x. If we let $\chi_{I,U}$ be the indicator function for this set, then $\chi_{I,U}(\mathbf{T} - x)$ as x ranges through some subset of $\mathbb{R}^d$ will count the number of times a patch from I appears in $\mathbf{T}$ in that subset, without overcounting.

Proposition 3.4. *Cylinder sets given by trim sets generate the metric topology on* Ω.

Proof. We establish that every ball of radius ϵ around a tiling $\mathbf{T}$ can be obtained as a cylinder set. Take the smallest patch in $\mathbf{T}$ that contains $B_{1/\epsilon}(0)$ and call it P, and denote by x the control point of a tile in P containing the origin. The set of all patches that are within ϵ of P can be partitioned into a trim set of translation classes I: take all patches P' that have a control point at x and for which $d(P, P') < \epsilon$. Then $\Omega_{I,B_{1/\epsilon}(0)}$ is a cylinder set that equals the ball of radius ϵ around $\mathbf{T}$. □

3.3. Translation-invariant measures and patch frequency

We begin this discussion by reviewing how translation-invariant Borel probability measures can be used to compute frequencies in the FLC case. Given some finite patch P, if U is a sufficiently small open set and μ is an invariant measure we can define the frequency of P to be $\text{freq}_\mu(P) = \dfrac{\mu(\Omega_{P,U})}{\text{Vol}(U)}$. If μ is ergodic then by the ergodic theorem for μ-a.e $\mathbf{T}$ we have

$$\text{freq}_\mu(P) = \lim_{R\to\infty} \frac{1}{\text{Vol}(B_R(0))\,\text{Vol}(U)} \int_{B_R(0)} \chi_{P,U}(\mathbf{T} - x)dx,$$

where $\chi_{P,U}$ is the indicator function for $\Omega_{P,U}$. The integral represents the number of times we see a copy of P in the ball of radius R around the origin in $\mathbf{T}$, so averaging this by the size of the ball gives us the frequency of P.

In the ILC case, when μ is a translation-invariant measure and I is a trim set we still see that $\mu(\Omega_{I,U})$ is a multiple of $\text{Vol}(U)$ for all sufficiently small sets U. Thus we can define the frequency of I to be $\text{freq}_\mu(I) = \dfrac{\mu(\Omega_{I,U})}{\text{Vol}(U)}$ as before. And as before we are justified in the use of the word "frequency" by the ergodic theorem. If I is a trim ϵ-ball around some patch P, then $\text{freq}(I)$ is the percent of time we see patches that look almost exactly like P.

Let $\mathcal{P}_n$ be the set of all connected n-tile patches that have a control point at the origin and are translates of patches that appear in Ω. The metric on patches gives us a measurable structure on $\mathcal{P}_n$, and since every subset of $\mathcal{P}_n$ is trim, freq_μ

forms a measure on $\mathcal{P}_n$. If we want, we can consider $\mathcal{P}_\infty = \bigcup \mathcal{P}_n$, which is not itself a trim set. However, since any subset of a trim set is trim, we can consider freq_μ to be a measure on (the set of measurable subsets of) any trim subset of $\mathcal{P}_\infty$.

If μ is a probability measure, then the frequency measure on $\mathcal{P}_n$ is *volume-normalized*, meaning that $\int_{\mathcal{P}_n} \text{Vol}(P)\, \text{freq}_\mu(dP) = 1$. This follows from the fact that $\mu(\Omega) = 1$ and Ω can be arbitrarily finely approximated by cylinder sets of the form $\Omega_{I_n(j_n),U_n}$, where $I_n(j_n) = B_\epsilon(P_n(j_n))$ and $U_n = \text{supp}(P_n(j_n))$ for some representative set of n-supertiles.

3.4. Entropy and complexity

We develop a notion of complexity based on the standard form in symbolic dynamics, but taking ideas from topological pressure theory and the topological entropy of flows. The complexity function distinguishes the sort of infinite local complexity represented by the solenoid (Example 7) from that of, say, tilings which have a higher topological dimension than their ambient dimension (Example 6, for instance).

There are three interrelated ways to define the complexity function, all of which yield slightly different actual numbers but have the same asymptotics and are based on the idea that complexity should count the number of patches of size L one might see in Ω. To that end we define a metric d_L on Ω for each $L > 0$ by

$$d_L(\mathbf{T}, \mathbf{T}') = \sup_{x \in [0,L]^d} \{d(\mathbf{T} - x, \mathbf{T}' - x)\}.$$

Two tilings will be within ϵ of one another in this d_L measure if their patches on $[-1/\epsilon, L + 1/\epsilon]^d$ are within ϵ in the patch metric. Our complexity functions will count up how many such patches there are.

For any $\epsilon > 0$ and $L > 0$ we define $N_1(\epsilon, L)$ to be the minimum number of balls of d_L-radius ϵ it takes to cover Ω. We define $N_2(\epsilon, L)$ to be the minimum number of sets of d_L-diameter ϵ it takes to cover Ω. It is clear that since every open cover using balls of d_L-radius ϵ is a cover by sets of d_L-diameter 2ϵ, we know that $N_2(2\epsilon, L) \leq N_1(\epsilon, L)$.

Our third version of a complexity function relies on the idea of an *ϵ-separated set:* a set of tilings in Ω, no two of which are within ϵ of each other in the d_L metric. We define $N_3(\epsilon, L)$ to be the maximum cardinality of an ϵ-separated set. If we have such a set then we can cover Ω with balls of d_L-radius ϵ centered on its elements, so we have that $N_3(\epsilon, L) \geq N_1(\epsilon, L)$. Also, since any set of diameter ϵ can contain at most one element of an ϵ-separated set, we have that $N_2(\epsilon, L) \geq N_3(\epsilon, L)$. Thus we have:

$$N_2(2\epsilon, L) \leq N_1(\epsilon, L) \leq N_3(\epsilon, L) \leq N_2(\epsilon, L).$$

If we let N denote any of these complexity functions, we can look at what happens as ϵ goes to 0 and/or as $L \to \infty$. For any given L we see that even for tilings with finite local complexity $\lim_{\epsilon \to 0} N(\epsilon, L) = \infty$. Instead we should fix an ϵ and investigate $\lim_{L \to \infty} N(\epsilon, L)$. We say that Ω has *bounded complexity* if $N(\epsilon, L)$

is bounded by some function of ϵ, independent of L. We say it has *polynomial complexity* if $N(\epsilon, L)$ is bounded by $C(\epsilon)(1+L)^{\alpha}$, where C is some function of ϵ and α is some positive constant.

Definition 3.5. The ϵ-*entropy* of the tiling dynamical system $(\Omega, \mathbb{R}^d)$ is given by

$$h_\epsilon(\Omega) = \limsup_{L\to\infty}(log(N(\epsilon, L)))/L^d.$$

If $\lim_{\epsilon\to 0} h_\epsilon(\Omega) = h(\Omega)$ is finite, then we say the system has finite entropy equal to $h(\Omega)$.

The usual complexity function $c(n)$ for a one-dimensional symbolic sequence on a finite number of letters counts the number of distinct words of length n. If we consider the sequence to be a tiling with labelled unit interval tiles, then any of our complexity functions $N(\epsilon, L)$ are approximately equal to $c([L + 2/\epsilon])/\epsilon$.

4. Hierarchical tilings: substitution and fusion

An important theme in the study of aperiodic order is hierarchical structures: sequences or tilings that can be seen as possessing structure at arbitrarily large length scales. The earliest work in this direction was on substitution sequences, which are surveyed in [17]. Self-similar tilings were a natural generalization to the tiling situation, and have also been studied extensively in the FLC case ([3] is an excellent reference). However, such hierarchical construction methods can lead naturally to tilings with infinite local complexity. Early examples of tilings with infinite local complexity arose from tilings with a finite number of tile sizes and a substitution algorithm that forced the tiles to slide past one another in infinitely many ways [6, 16]. We have selected three examples that show some of the things that can happen when infinite local complexity arises in a hierarchical tiling.

4.1. Generating hierarchical tilings I: substitution

The earliest form of substitution was for symbolic systems, where there is some discrete alphabet $\mathcal{A}$ and some substitution rule $\sigma : \mathcal{A} \to \mathcal{A}^*$ that takes letters to words. For instance, the Fibonacci substitution has $\mathcal{A} = \{a, b\}$, with $\sigma(a) = ab$ and $\sigma(b) = a$. One can iterate the substitution by substituting each letter individually and concatenating the results. In the Fibonacci example we have

$$\sigma^2(a) = \sigma(a)\sigma(b) = ab\,a \qquad \sigma^3(a) = \sigma(a)\sigma(b)\sigma(a) = ab\,a\,ab$$

and so on. One can generate infinite sequences in this manner.

Extending this to the tiling case in one dimension is simple because tiles are intervals and can be concatenated without discrepancy. However, once we are in two dimensions the tiles have geometry that can prevent the tiles from fitting together. The first way around this was to devise inflate-and-subdivide rules that generate self-similar or self-affine tilings via linear expanding maps. Many beautiful examples have been discovered and investigated, and can be found on the Tilings Encyclopedia website [13].

An inflate-and-subdivide rule requires a linear expansion map $\phi : \mathbb{R}^d \to \mathbb{R}^d$ such that for each prototile $t \in \mathcal{P}$, the expanded set $\phi(\text{supp}(t))$ can be expressed as a union of tiles equivalent to prototiles from $\mathcal{P}$. We write $S(t)$ to represent the patch of tiles that result from the inflate-and-subdivide process, called a 1-supertile. We can apply the substitution rule to the tiles in the patch $S(t)$ to obtain the patch $S^2(t)$, which we call a 2-supertile. Repeated substitution produces higher-order supertiles that grow to cover $\mathbb{R}^d$ in the limit.

Example 5. *Infinitely many tile lengths.* Let $[1, 3]$ be both the label set and the set of tile lengths for a one-dimensional tiling as in our first introductory example. For $x \in [1, 3]$, the tile denoted t_x is taken to be of type x and $\text{supp}(t_x)$ is an interval of length x. The control points are taken to be the left endpoints. We can take the metric on the label set to be given by $d_{\mathcal{L}}(x, y) = |x - y|/2$, which is somewhat arbitrary but agrees with the Hausdorff distance of two tiles of lengths x and y that have the same midpoint.

We define a substitution rule that inflates by the expansion map $\phi(x) = 3x/2$ and subdivides the result only if it is larger than 3. If $x \in [1, 2]$ we define $S(t_x) = t_{3x/2}$, supported in the interval $\phi(\text{supp}(t_x))$. If $x \in (2, 3]$ we define $S(t_x) = t_x \cup t_{x/2}$, again supported in the interval $\phi(\text{supp}(t_x))$. Notice that the substitution rule is discontinuous: two tiles with lengths on either side of 2 substitute to patches that are not close in the patch metric since they have different numbers of tiles.

In Figure 1 we show a 9-supertile for the substitution, with the interval lengths coded by color [greyscale]; tiles close in color are close in length. We see consecutive copies of the same tile on three occasions.

FIGURE 1. Nine iterations of the substitution rule, applied to $\pi/2$.

A version of the tiling space generated by this rule, considered from a fusion standpoint, is studied in [12], where it is shown to be minimal and have a totally disconnected transversal and a unique translation-invariant Borel probability measure that is nonatomic. We investigate more about this tiling space in Section 6.1.

The geometric rigidity imposed by the linear map ϕ can be loosened somewhat. Tiling substitution rules exist such that any tile t is substituted by a patch of tiles $S(t)$, but this patch may not be supported on a set that is a linear expansion of t. These have been called combinatorial substitutions [8], a special case of which is known by the term "generalized substitutions" [2].

A straightforward way to generate tiling substitutions in $\mathbb{R}^d$ is to begin with the direct product of d one-dimensional substitutions. Given d substitutions $\sigma_1, \sigma_2, \ldots, \sigma_d$ on alphabets $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_d$ we can define

$$\sigma(a_1, a_2, \ldots, a_d) = (\sigma_1(a_1), \sigma_2(a_2), \ldots, \sigma_d(a_d))$$

A tile associated with the label $(a_1, a_2, \ldots, a_d)$ is a d-dimensional rectangle, the length of the ith side depending on $a_i \in \mathcal{A}_i$. In a direct product tiling substitution the tiles must line up facet-to-facet and thus always have finite local complexity.

These can be made into the more interesting *direct product variation* (*DPV*) substitutions, one of which is the example below. To construct such a substitution we rearrange the inside of at least one of the substituted tiles in order to break the direct product structure. Care must be taken to ensure that the rearranged interior still forms a legal patch when substituted so that the substitution admits tilings.

Whether there is finite or infinite local complexity depends on combinatorial, number-theoretic, and/or geometric details. One with ILC, based on the product of $a \to abbb, b \to a$ with itself, is the primary ILC example in [9] and requires four tile sizes. The simpler example we present here is similar to the one whose cohomology was computed in [10].

Example 6. *Direct product variation (DPV).* Let $\sigma_1 : a \to abbb, b \to a$ and $\sigma_2 : c \to cc$. There are two rectangular tile types we call $A = a \times c$ and $B = b \times c$, where we think of a, b, and c as representing both intervals, their lengths, and their labels. The direct product will then be

$$\boxed{A} \to \begin{array}{|c|c|c|c|}\hline A & B & B & B \\ \hline A & B & B & B \\ \hline \end{array}, \qquad \boxed{B} \to \begin{array}{|c|}\hline A \\ \hline A \\ \hline \end{array}$$

We can vary this direct product as follows and be guaranteed that the substituted tiles will still fit together.

$$\boxed{A} \to \begin{array}{|c|c|c|c|}\hline A & B & B & B \\ \hline B & B & B & A \\ \hline \end{array}, \qquad \boxed{B} \to \begin{array}{|c|}\hline A \\ \hline A \\ \hline \end{array}$$

By varying the widths of A and B we can obtain tilings with either finite or infinite local complexity. If the widths are irrationally related, the substitution rule admits tilings with ILC. Figure 2 shows three iterations of the A tile using the widths $a = (1 + \sqrt{13})/2, b = 1$, which are the natural widths for the self-affine tiling for this substitution. Horizontal fault lines are beginning to develop, with mismatches between the tiles above the lines and those below. Each iteration of the substitution produces new offsets along the fault lines, ultimately resulting in infinite local complexity. The connection between these fault lines and the projection method will be discussed when we fully analyze this tiling in Section 6.2.

4.2. Generating hierarchical structures II: fusion

Like substitution, fusion constructs tilings by making a series of n-supertiles that get larger and larger at each level. The difference is that while substitution constructs an n-supertile by replacing each tile in an $(n-1)$-supertile with a substituted tile, fusion constructs an n-supertile by concatenating or 'fusing' a number of $(n-1)$-supertiles. We could think of substitution as being a cellular model: each tile is a cell that can expand and subdivide itself into new cells. Fusion is an atomic

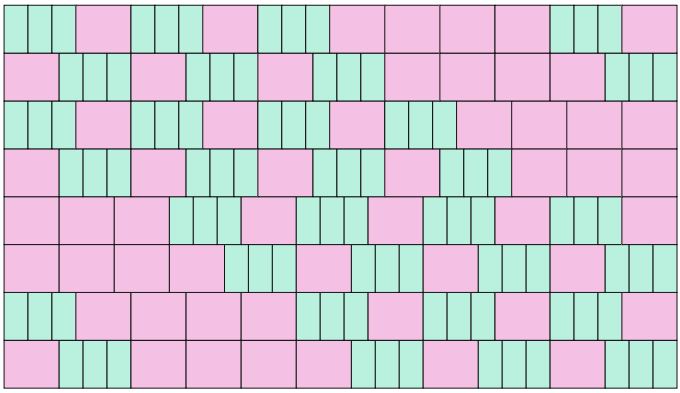

FIGURE 2. A direct product variation 3-supertile

model: each tile is an atom that can bond to other atoms to form molecules, which can themselves bond together to form larger structures. We refer the reader to [11, 12] for technical details and more examples, but we describe many key points here.

The prototile set $\mathcal{P}_0$ will serve as our 0-supertiles. Recall that there is a compact label set which we now call $\mathcal{L}_0$ that labels the prototiles and generates a tile and patch metric. The 1-supertiles are defined to be a set of finite patches $\mathcal{P}_1$ of tiles from $\mathcal{P}$. We require that there be a compact $\mathcal{L}_1$ that labels the 1-supertiles, so that we may write $\mathcal{P}_1 = \{P_1(c) \mid c \in \mathcal{L}_1\}$. It is convenient but not necessary to require that if $c_n \to c$ in $\mathcal{L}_1$, then $P_1(c_n) \to P_1(c)$ in the patch metric generated by $\mathcal{L}_0$. There are examples where the fusion and/or substitution is only piecewise continuous, for instance Example 5.

We make our set of 2-supertiles $\mathcal{P}_2$ by requiring that each element of $\mathcal{P}_2$ be a *fusion* of 1-supertiles: a finite, connected union of patches that overlap only on their boundaries. We require that $\mathcal{P}_2$ is labelled by some compact label set $\mathcal{L}_2$, and we write $\mathcal{P}_2 = \{P_2(c) \mid c \in \mathcal{L}_2\}$. It is convenient if the patch metric generated by $\mathcal{L}_2$ is compatible with the patch metrics generated by $\mathcal{L}_1$ and $\mathcal{L}$ in the sense of the previous paragraph.

We continue in this fashion, constructing our n-supertiles as fusions of $(n-1)$-supertiles and requiring that each $\mathcal{P}_n$ be labelled by a compact set $\mathcal{L}_n$. The *fusion rule* $\mathcal{R}$ is the set of all supertiles from all levels:

$$\mathcal{R} = \{\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2, \dots\}$$

We say a tiling $\mathbf{T}$ is *admitted* by $\mathcal{R}$ if every patch of tiles in $\mathbf{T}$ is equivalent to one appearing inside of a supertile from $\mathcal{R}$. The tiling space $\Omega_{\mathcal{R}}$ is the set of all tilings admitted by $\mathcal{R}$, and $\mathcal{R}$ can be thought of as the language of $\Omega_{\mathcal{R}}$. In order for a fusion rule to admit an infinite tiling, the sizes of supertiles must be unbounded. Moreover, $\Omega_{\mathcal{R}}$ is a translation-invariant tiling space and can be analyzed dynamically.

In order to avoid trivialities, we assume that $\Omega_{\mathcal{R}}$ is not empty and that every element of $\mathcal{P}_n$, for each n, appears somewhere inside an infinite tiling in $\Omega_{\mathcal{R}}$. We can assume the latter without loss of generality since any superflous supertiles can be removed from $\mathcal{R}$ without changing $\Omega_{\mathcal{R}}$.

Definition 4.1. An *infinite-order supertile* P_∞ is a tiling of an unbounded region of $\mathbb{R}^d$ for which there is a sequence of supertiles $P_n \subset \mathcal{P}_n$ and translations $x_n \in \mathbb{R}^d$ for which $P_\infty = \lim_{n\to\infty}(P_n - x_n)$ and $P_n - x_n \subset P_{n+1} - x_{n+1}$ for all n.

Tilings admitted by a FLC fusion rule $\mathcal{R}$ are either one or the concatenation of finitely many infinite-order supertiles. In the ILC situation there is another possibility, that they are the limit of such infinite-order supertiles in the big ball metric.

Example 7. *Solenoid extensions*[1]. We present a simple family of fusions that are not substitutions. All of the examples are measurably conjugate to the dyadic solenoid system, which is described as an inverse limit in Chapter 3 of this volume (pp. 73–104) and can be seen as a height-1 suspension of the dyadic odometer. However, the topology is highly sensitive to changes in the prototile set.

For this family of tilings, the support of all prototiles is $[0, 1]$. The label set $\mathcal{L}$ is a compactification of the non-negative integers $\mathbb{N}_0$ and we write $\mathcal{L} = \mathbb{N}_0 \cup \mathcal{L}_c$. We denote a tile of type l as A_l. Regardless of the specific nature of $\mathcal{L}$ the fusion rule will be constructed as follows. Letting the set of 0-supertiles, $\mathcal{P}_0$ be the prototile set, we construct our set of 1-supertiles as follows. For $l \in \mathcal{L}_c$ and $l = 1, 2, 3, 4, \dots$ we define $P_1(l) = A_l \cup (A_0 + 1)$, that is, the two-tile patch supported on $[0, 2]$ given by the concatenation of A_l and A_0. We hope it does not risk too much confusion about the precise support of $P_1(l)$ if we abuse notation and write $P_1(l) = A_l A_0$. We write $\mathcal{P}_1 = \{P_1(l), l \in \mathcal{L}_1\}$, where $\mathcal{L}_1$ is the subset of $\mathcal{L}$ given by $\{1, 2, 3, 4, \dots\} \cup \mathcal{L}_c$. Now to generate the 2-supertiles we concatenate each 1-supertile with label in $\{2, 3, 4, \dots\} \cup \mathcal{L}_c$ with $P_1(1)$, so that we have $P_2(l) = P_1(l)P_1(1) = A_l A_0 A_1 A_0$. (Again we abuse notation but we know this is supported on $[0, 4]$.) The set of 2-supertiles takes the form $\mathcal{P}_2 = \{P_2(l), l \in \mathcal{L}_2\}$, where $\mathcal{L}_2$ is the subset of $\mathcal{L}$ given by $\{2, 3, 4, \dots\} \cup \mathcal{L}_c$. Similarly the set of 3-supertiles will have the form $P_3(l) = P_2(l)P_2(2) = A_l A_0 A_1 A_0 A_2 A_0 A_1 A_0$, for $l \in \mathcal{L}_3$. The general form for the set of k-supertiles, $k = 1, 2, 3, \dots$ is

$$\mathcal{P}_k = \{P_{k-1}(l)P_{k-1}(k-1), l \in \mathcal{L}_k\}.$$

By looking at the form for $P_3(l)$, we see that there will be an A_0 in every other slot, an A_1 in every fourth slot, and can surmise that there will be an A_n in every 2^{n+1}th slot. In fact we can generate an infinite tiling admitted by this fusion rule by a method quite similar to the construction of a Toeplitz sequence. We begin by placing infinitely many A_0's on the line with a unit space between them. Of the remaining spaces, we alternate by filling one with A_1 and leaving the next one empty. We continue in this fashion, filling every other of the remaining

[1] These are closely related to discrete actions known as "Toeplitz flows" and are surveyed in [7].

spaces with an A_2, and so on. When the process is finished, there may or may not be one empty space. If there is, it should be filled with an A_l with $l \in \mathcal{L}_c$.

Whatever form $\mathcal{L}_c$ takes, the tiling space admitted by this fusion rule has infinite local complexity. When $\mathcal{L}_c$ consists of a single point we will see in Section 6.3.5 it is actually less 'complex' than FLC examples because its complexity is bounded. When $\mathcal{L}_c$ is finite but greater than 1 we can use Theorem 5.9 to show that it is topologically conjugate to a tiling space with finite local complexity. We will look closely a few special cases in Section 6.3.

Tilings generated by substitution can always be seen as being generated by fusion since a tile that has been substituted n times can be seen as the concatenation of tiles that have been substituted $(n-1)$ times. However the converse is not true. Fusion is more general and can allow us to vary the size of the supertile sets from one level to the next. They also can allow us to vary the fusion patterns from level to level, and can account for generalized substitutions and random substitutions.

4.3. Transition matrices and transition maps

4.3.1. Transition matrices. When there is an inflate-and-subdivide rule or a substitution rule for a tiling with a finite number of tile types it is very handy to compute the *transition matrix* A of the substitution. The (i, j)th entry of the matrix is given by the number of tiles of type i in the substitution of the tile of type j. The matrix A^n knows how many of each prototile type can be found in the n-supertiles of each type. We say A is *primitive* if there is some n for which A^n has all positive entries. This means that each n-supertile contains copies of every tile type.

In the planar case, when the inflate-and-subdivide rule is a similarity, we can consider the expansion to be by some complex number λ, which is the Perron eigenvalue of the transition matrix. Early in the study of self-similar tiling dynamical systems it was discovered that the algebraic type of λ had a significant impact on the dynamics. Working in the context of finite local complexity, Solomyak [22] showed that under the conditions of primitivity and recognizability a self-similar tiling of the line fails to be weakly mixing if and only if $|\lambda|$ is a Pisot number (an algebraic integer, all of whose algebraic conjugates are less that one in modulus). In the same work he showed that a tiling of the (complex) plane fails to be weakly mixing if and only if λ is a complex Pisot number.

It turns out that the algebraic type of the expansion constant also has an effect on local complexity. In [9] it is shown that under common conditions, if the length expansion is a Pisot number then tilings admitted by the substitution must have finite local complexity. In this situation it is known that there must be some measurable spectrum and cannot be weakly mixing [22]. Thus weak mixing and local complexity are linked via the expansion constant; when it is not Pisot there is a chance of the local complexity becoming infinite.

If a fusion rule has a finite number of n-supertiles at each stage, we only need to generalize the idea of a transition matrix for substitutions slightly. We now have transition matrices $A_{n,N}$ whose (i, j)th entry represents the number of n-supertiles

of type i in $P_N(j)$. For any m between n and N we have $A_{n,N} = A_{n,m}A_{m,N}$. Thus even though we are unable to use Perron–Frobenius theory, we are able to use a parallel type of analysis to determine the possible invariant measures and patch frequencies. If the tiling space has ILC as in Example 6, the patches that only appear along infinite fault lines have frequency 0 and all other patches have nonzero frequency [12].

4.3.2. Transition maps. When we have a substitution or fusion on infinitely many prototiles or supertiles, the transition between levels must be a map $A_{n,N} : \mathcal{P}_n \times \mathcal{P}_N \to \mathbb{Z}$ for which $A_{n,N}(P, Q)$ represents the number of n-supertiles of type P in the N-supertile of type Q. For fixed Q, there are only finitely many nonzero entries, but a fixed $P \in \mathcal{P}_n$ might appear in infinitely many tiles. For instance, every tile type in the solenoid example has this property except ones with label in the compactification $\mathcal{L}_c$.

Even though the transition maps are no longer matrices, they can still be used to obtain the possible invariant measures for a large class of ILC fusions (see Section 5). The overarching principle is that we have a measure ρ_n on each supertile set $\mathcal{P}_n$ so that for any measurable, trim subset $I \subset \mathcal{P}_n$, $\rho_n(I)$ represents the frequency of seeing any supertile from I in a tiling. When the sequence of such frequencies $\{\rho_n\}$ behave nicely with respect to transition we obtain both a translation-invariant measure on the tiling space and a handy formula for computing the frequencies of all types of patches, not just supertiles.

4.4. Recognizable, van Hove, and primitive fusion rules

The definition of a fusion rule is sufficiently general as to encompass all tilings whatsoever, and therefore we need to put some restrictions on the rules to make them meaningful. The three standard assumptions are that the fusion rule be recognizable, van Hove, and primitive. Recognizability and the van Hove property are defined for fusions the same way whether the local complexity is finite or infinite, but primitivity requires a more subtle definition in the case of infinite local complexity.

Recognizability means that the substitution or fusion can be undone in a well-defined way. For self-similar tilings, recognizability means that there is some finite 'recognizability radius' R such that if $\mathbf{T}$ and $\mathbf{T}'$ have identical patches in the ball $B_R(x)$, then they have the same substituted tile at x, situated in precisely the same way. For fusion rules, we need a recognizability radius for each level of supertiles. That is, for each $n \geq 1$ there is an $R_n > 0$ such that if $\mathbf{T}$ and $\mathbf{T}'$ have the same patch of $(n-1)$-supertiles in $B_{R_n}(x)$, then they have the same n-supertile at x in precisely the same location and orientation.

If Ω is either a substitution or fusion tiling space, we can define spaces Ω^n to be the space of tilings from Ω with the n-supertiles considered to be the set of prototiles by 'forgetting' all the tiles in their interiors. Since every tiling in Ω is a union of n supertiles for any n, this is a well-defined tiling space. There is always a map from Ω^n to Ω^{n-1} since we know how each n-supertile is constructed from

$(n-1)$-supertiles. However, this map is not necessarily invertible: there could be tilings in Ω^{n-1} that could be composed into tilings of n-supertiles in more than one way. The substitution or fusion rule is recognizable if that is not the case and the map is a homeomorphism for each n.

It is convenient to work with fusion rules for which all of the supertiles grow in area in a reasonable way, for instance without becoming arbitrarily long and skinny. One way to avoid this is to require that the boundaries of supertiles are small relative to their interiors. To this end, for $r > 0$ and any set $U \in \mathbb{R}^d$ we define $(\partial(U))^{+r}$ to be the set of all points in $\mathbb{R}^d$ that are within r of the boundary of U. A sequence of sets $\{U_n\}$ in $\mathbb{R}^d$ is called *van Hove* if for every $r \geq 0$ we have that $\lim_{n\to\infty} \dfrac{\mathrm{Vol}(\partial(U_n)^{+r})}{\mathrm{Vol}(U_n)} = 0$. A fusion rule is van Hove if any sequence of n-supertiles $\{P_n\}$, where $P_n \in \mathcal{P}_n$, is supported on a van Hove sequence. This property is sufficient to ensure that the tiling space is not empty, but it is not necessary.

The general idea behind primitivity is that given any n, there should be some N for which every N-supertile contains n-supertiles of each type. This definition makes sense when the sets of supertiles are finite. When they are not, we simply require that for any n and any open set I of n-supertiles, there is an N such that every N-supertile contains an n-supertile from I. Thus, in the ILC case is that the size of I affects the size of N, whereas in the FLC case a single N can be chosen for all n-supertiles. A primitive fusion or substitution has the property that for sufficiently large N, $A_{n,N}(I, Q) \neq 0$ for all $Q \in \mathcal{P}_N$.

5. Results about ILC tilings

5.1. Fault lines and fractured tiling spaces

We have seen that it is easy to construct tilings of the plane with infinite local complexity. An important question to ask is, what are the ways a planar tiling can have ILC? When the prototile set is finite, the answer is given in [16]. Up to translation it can only happen if there are arbitrarily long line segments composed of tile edges where the tiles meet up in arbitrarily many different ways.

Theorem 5.1 ([16]). *A tiling of the plane with translated copies of a finite set of tiles either has only a finite number of local configurations or else contains arbitrarily long line segments in the boundaries of the tiles.*

Kenyon also shows in [16] that if we allow infinitely many rotations of our finite prototile set there is only one additional way that infinite local complexity can arise: Some of the tile edges would have to have circular arcs, so that a circular patch of tiles can be constructed. This patch of tiles could then be rotated by arbitrary amounts inside any patch of tiles that surrounds it to create infinitely many different patches.

Kenyon called an infinite line of tile boundaries along which tiles can slide an *earthquake* and elsewhere in the literature it is often called a *fault line*. Unfortunately fault lines do not have a unified definition, so we provide one here.

Definition 5.2. A tiling $\mathbf{T}$ in a tiling space Ω is said to have a *fault line* ℓ if there are infinitely many nonequivalent tilings $\mathbf{T}'$ such that $\mathbf{T}' = \mathbf{T}$ on one side of ℓ and $\mathbf{T}' = \mathbf{T} - x$ on the other side of ℓ for some $x \in \mathbb{R}^2$ that is parallel to ℓ.

However, fault lines are not topologically invariant: one can take a tiling space containing fault lines and relabel all the tiles by 'collaring' (see Chapter 3 of this book, pp. 73–104): each tile is labelled by its corona. The resulting tiling space will not contain fault lines per se, but they will still be *fractured*. We offer here a new definition that is not specific to planar tilings and may be related to the proximal and asymptotic structure of tiling spaces (see [4]).

Definition 5.3. A space Ω of tilings of $\mathbb{R}^d$ has a *fracture in the direction of* $x \in \mathbb{R}^d$ if there exists some $y \in \mathbb{R}^d$ and two tilings $\mathbf{T}, \mathbf{T}' \in \Omega$ such that

$$\lim_{t\to\infty} d(\mathbf{T} - ty, \mathbf{T}' - ty) = 0 \quad \text{and} \quad \lim_{t\to-\infty} d(\mathbf{T} - ty, \mathbf{T}' - x - ty) = 0.$$

So a tiling space is fractured in the direction x if there are tilings that asymptotically agree in the y direction and, after an offset, asymptotically agree in the $-y$ direction. There can be a large region 'in the x direction' in the middle of the tiling on which they do not agree. In one dimension, a tile near the origin could be added, removed or resized. In two dimensions, tilings with fault lines are fractured, and so are any tilings that are MLD to them. Moreover tilings that are asymptotically proximal will also have fractured tiling spaces. As an example, consider a chair tiling with an infinite diagonal of chairs that can be completely flipped without altering the rest of the tiling.

Notice that y is not uniquely defined. It obviously can be rescaled, but in two or higher dimensions the direction can be changed. The direction of x, however, cannot be changed without changing the direction of the fracture. A tiling space can have multiple fractures in different directions: see [8] for a planar tiling with translationally-finite prototile set that has fault lines in three independent directions.

Fractures and fault lines play an important role in the spectrum of tiling dynamical systems. Recall the following definition from higher-dimensional dynamics.

Definition 5.4. The dynamical system $(\Omega, \mathbb{R}^d)$ has an eigenfunction $f : \Omega \to \mathbb{C}$ with eigenvalue $\alpha \in \mathbb{R}^d$ if for any $\mathbf{T} \in \Omega$ and $y \in \mathbb{R}^d$,

$$f(\mathbf{T} - y) = \exp(2\pi i \alpha \cdot y) f(\mathbf{T}).$$

Theorem 5.5. [2] *If $f \in C[\Omega, \mathbb{C}]$ is a continuous eigenfunction with eigenvalue $\alpha \in \mathbb{R}^d$ and Ω has a fracture in the direction of x, then $\alpha \cdot x$ is an integer.*

Proof. Suppose that $y \in \mathbb{R}^d$ satisfies the fracture definition for tilings $\mathbf{T}, \mathbf{T}' \in \Omega$. Since f is uniformly continuous, $0 = \lim_{t\to\infty} (f(T - ty) - f(T' - ty)) = \lim \exp(2\pi i t\alpha \cdot y)[f(T) - f(T')]$, so $f(T) = f(T')$. Similarly, taking a limit as $t \to -\infty$, we

[2] This result is joint with L. Sadun.

get $f(T) = f(T' - x)$. But then $f(T') = f(T' - x) = \exp(2\pi i\alpha \cdot x)f(T')$, so $\exp(2\pi i\alpha \cdot x) = 1$ and the result is proved. □

Corollary 5.6. *If the tiling space contains fractures in the direction of kx for multiple values of $k \in \mathbb{R}$, then $k\alpha \cdot x \in \mathbb{Z}$ for each of them. If any of the k's are irrationally related, or if they can be arbitrarily small, then $\alpha \cdot x = 0$. If the k's have a greatest common factor m, then $\alpha \cdot x \in \frac{1}{m}\mathbb{Z}$.*

We will use this corollary to compute the topological spectrum of certain direct product variation substitutions in Section 6.2.

5.2. General results for ILC fusion tilings

If a fusion or substitution rule is primitive, then the tiling or sequence system associated with it is necessarily minimal regardless of local complexity [11, 12]. It is possible for a fusion tiling space or substitution sequence space to be minimal even if it is not primitive, the Chacon substitution and its DPV analogues being examples. If a fusion is both recognizable and van Hove, primitivity and minimality are equivalent:

Theorem 5.7 ([12]). *If the fusion rule $\mathcal{R}$ is primitive, then the fusion tiling dynamical system $(\Omega_{\mathcal{R}}, \mathbb{R}^d)$ is minimal. Conversely, a recognizable, van Hove fusion rule that is not primitive cannot have a minimal dynamical system.*

In Section 3.2 we defined a trim set of patches (basically, one that contains no repeats up to translation), and we explained in Section 3.3 why it is appropriate to consider $\text{freq}_\mu(P) = \dfrac{\mu(\Omega_{I,U})}{\text{Vol}(U)}$ to be the frequency of occurrence of patches from I throughout Ω (basically, the ergodic theorem). For fusion tilings, we can say more about frequencies of patches by understanding frequencies of supertiles. While our discussion is written in the context of ILC fusions, we note that the construction can be simplified to apply to the FLC case.

Throughout this discussion it is essential that our fusion rule be recognizable and van Hove, and we refer the reader to [12] for proofs and more details. Consider a set $I \subset \mathcal{P}_n$ of n-supertiles. Supposing that we have chosen the control points of our n-supertiles so that each supertile contains some ϵ-ball around the origin, I is automatically a trim set since by recognizability a tiling cannot have more than one supertile at any given location. In a slight abuse of notation we define $\Omega_{I,U}$ to be the set of all tilings that contain an n-supertile from I at a location determined by U[3].

We define a measure ρ_n on $\mathcal{P}_n$ defined by $\rho_n(I) = \dfrac{\mu(\Omega_{I,U})}{\text{Vol}(U)}$. When μ is a probability measure, each ρ_n is *volume normalized* in that $\displaystyle\int_{P\in\mathcal{P}_n} \text{Vol}(P)d\rho_n = 1$.

[3]The set I of n-supertiles corresponds, by recognizability, to a trim set of patches of ordinary tiles, each of which is larger than its supertile by some recognizability radius.

This can be seen as follows. Suppose that $I_1, I_2, \ldots, I_j$ is a partition of $\mathcal{P}_n$ into sets of very small diameter and that $P_k \in I_k$ for all $k = 1, \ldots, j$. Then $\bigcup_{P_k} \Omega_{P_k, \mathrm{Vol}(P_k)} \approx \Omega$ and so $\sum_{k=1}^{j} \rho_n(P_k) \mathrm{Vol}(P_k) \approx \mu(\Omega) = 1$.

The transition map $A_{n,N}$ can be thought of as inducing a map from measures on $\mathcal{P}_N$ to measures on $\mathcal{P}_n$. For a fixed trim set $I \subset \mathcal{P}_n$, the function $A_{n,N}(I, Q)$ can be defined as the number of times an n-supertile from I is contained in the N-supertile Q. Since each such Q contains only finitely many n-supertiles this function takes values in the nonnegative integers.

Let $n < N$ be fixed, let ν_N be a measure on $\mathcal{P}_N$ and let $I \subset \mathcal{P}_n$ be a trim set. We define a measure ν_n on $\mathcal{P}_n$ as

$$\nu_n(I) = (A_{n,N}\nu_N)(I) = \int_{Q \in \mathcal{P}_N} A_{n,N}(I, Q) d\nu_N.$$

We say that a sequence of measures $\{\nu_n\}_{n=0}^{\infty}$ is *transition-consistent* if whenever $n < N$, $\nu_n = A_{n,N}\nu_N$. This means that the frequencies that ν_n assigns to n-supertiles is consistent with the frequencies that ν_N assigns to N-supertiles, for any $n \leq N$.

The fact that translation-invariant probability measures give rise to sequences of volume normalized and transition consistent supertile measures and vice versa is the subject of the next theorem.

Theorem 5.8 ([12]). *Let $\mathcal{R}$ be a fusion rule that is van Hove and recognizable. Each translation-invariant Borel probability measure μ on $\Omega_{\mathcal{R}}$ gives rise to a sequence of volume normalized and transition consistent measures $\{\rho_n\}$ on $\mathcal{P}_n$. Moreover, for any trim set of patches I*

$$\mathrm{freq}_{\mu}(I) = \lim_{n \to \infty} \int_{P \in \mathcal{P}_n} \#(I \text{ in } P) d\rho_n, \tag{4}$$

where $\#(I \text{ in } P)$ denotes the number of translates of patches in the family I that are subsets of P. Conversely, each sequence $\{\rho_n\}$ of volume normalized and transition consistent measures corresponds to exactly one invariant measure μ via equation (4).

This means that the invariant measures for a recognizable van Hove fusion system are almost completely determined by the transition maps, a fact that is consistent with substitution sequence and self-similar tiling theory. In those cases one looks at the transition matrix and uses the Perron–Frobenius theorem to conclude when the system is uniquely ergodic. In the FLC fusion case the transition maps are always matrices and we can parameterize the set of all possible measures by a Choquet simplex related to the matrix system.

5.3. ILC fusion tilings with finitely many n-supertiles

When there are finitely many prototiles and finitely many n-supertiles at every level, then any set of n-supertiles contains only finitely many patches of a given size. These patches are said to be *literally admitted* by the fusion rule. If the tiling space has infinite local complexity there can be patches that are not literally admitted; we call these *admitted in the limit*. An easy corollary to Theorem 5.8 is that the patches that are admitted in the limit have frequency 0. More precisely, if we take any trim set I of patches that are admitted in the limit, we clearly have that $\#(I \in P) = 0$ for any supertile P, so $\text{freq}_\mu(I) = 0$ by equation 4.

In fact, the set of literally admitted patches is countable, and they support the frequency measure. Thus the frequency measure is atomic and for any trim set of patches I, $\text{freq}_\mu(I) = \sum_{P \in I} \text{freq}_\mu(P)$. This is Theorem 4.4 of [12].

Often in the planar situation we can prove topological weak mixing or obtain strong restrictions on the topological spectrum by combining Theorem 5.1 and Corollary 5.6. The infinite local complexity comes from fault lines that in many examples have arbitrarily small shears. If this happens in two independent directions the tiling space in question is topologically weakly mixing.

5.4. Homeomorphism of ILC and FLC tiling spaces

Every non-periodic FLC tiling space is homeomorphic, or even topologically conjugate, to an ILC tiling space obtained by collaring with infinite collars. It is often possible to obtain an ILC tiling space from an FLC one geometrically, too. That makes it important to know whether a given space with infinite local complexity can be converted into one with finite local complexity. This classification problem is open in all dimensions except dimension one.

In one dimension it is possible to detect when a tiling space with ILC is homeomorphic to a tiling space with FLC. It is necessary but not sufficient for the tiling space to have a totally disconnected transversal and expansive dynamics since all tilings with FLC have those properties. However, it is possible to construct examples of ILC tilings that have these properties but cannot be homeomorphic to FLC tiling spaces, for instance by deforming the tiles of a solenoid extension [12]. The property a one-dimensional tiling space $(\Omega, \mathbb{R})$ needs is not expansivity, it is *strong expansivity*: the first return map to the transversal should be expansive as a $\mathbb{Z}$-action.

Theorem 5.9 ([12]). *If a one-dimensional tiling space is strongly expansive and the canonical transversal is totally disconnected, then it is homeomorphic to a one-dimensional tiling space with finite local complexity.*

There is a more general definition of strong expansivity for tilings of $\mathbb{R}^d$, but it seems unlikely that the d-dimensional analogue of Theorem 5.9 holds. A potential counterexample that is a variant of the pinwheel tilings appears in [12]. It has uncountably many ergodic measures and therefore is probably not homeomorphic to a tiling space with finite local complexity.

6. Analysis of our three main examples.

6.1. Example 5: variable size tile lengths

Consider Ω to be the tiling space admitted by the substitution in Example 5. Recall that the tile labels and lengths are in $[1, 3]$, the expansion factor is $3/2$, and an expanded tile is decomposed into two tiles of lengths $1/3$ and $2/3$ of the expanded tile if it is larger than 3 and otherwise not subdivided.

6.1.1. Minimality. The tiles in any tiling $\mathbf{T}$ all have lengths of the form $(2^j/3^k)y$, where $y \in [1, 3]$ and j and k are nonnegative integers. These lengths are dense in $[1, 3]$ for any fixed y, so we know that it is possible that the orbit of $\mathbf{T}$ is dense in Ω. We prove it by establishing primitivity.

The N-supertile of length L is made up of n-supertiles of lengths $(2^j/3^k)L$, for nonnegative integers $j \leq k$ in a range of values that depends more on N than it does on L. The largest value of k, denoted $\bar{k}$, comes from the rightmost n-supertile of L and has the property that $(2/3)^{\bar{k}}L \in [(3/2)^n, 3(3/2)^n]$. The smallest value of k, denoted $\underline{k}$, comes from the leftmost n-supertile and has the property that $(1/3)^{\underline{k}}L \in [(3/2)^n, 3(3/2)^n]$. All values of k between these two values occur, and all values of j for which $(2^j/3^k)L$ lie in range also appear, with the js being consecutive integers in the range.

Consider a set of n-supertiles of diameter ϵ, so that all of the lengths are within ϵ of each other. We can choose N such that for any length L, the maximum of k is sufficiently large so that every interval of size ϵ contains enough points of the form $(m/3^{\bar{k}})L$ that at least one of these points must be of the form $(2^j/3^k)L$, where $\underline{k} \leq k \leq \bar{k}$. This proves the substitution is primitive. Theorem 5.7 implies that the variable size substitution example must therefore have a minimal dynamical system.

6.1.2. Weak mixing. This system has no measurable or topological eigenvalues, as shown in this clever proof from L. Sadun. Let Ω be the tiling space from Example 5 and let Ω_λ be the tiling space obtained by expanding every tiling in Ω by the linear map $f(x) = \lambda x$ for some $\lambda > 1$. Then eigenvalues of Ω_λ are exactly the eigenvalues of Ω multiplied by $1/\lambda$.

On the other hand we can consider the tiles in Ω_λ, which live in $[\lambda, 3\lambda]$, to be lengths for supertiles in Ω, and apply the decomposition map for supertiles to each tiling in Ω_λ. This will result in tilings from Ω and we have a map from Ω_λ to Ω that is a bijection everywhere except on the set of measure 0 for which the subdivision rule is discontinuous. This means that Ω_λ and Ω are measurably conjugate and therefore have the same eigenvalues.

These two facts together mean that the set of eigenvalues of Ω (i.e., its spectrum) is invariant by scaling by any $\lambda > 1$. This means that the spectrum must either be $\mathbb{R}$, the nonnegative reals, the nonpositive reals, or 0. However, we know that since Ω is separable the spectrum must be countable. Thus the only possible measurable eigenvalue is 0. Since continuous eigenfunctions are measurable this also implies topological weak mixing.

6.1.3. Invariant measure. The analysis of the related fusion rule in [12] can be adapted to find the sequence $\{\rho_n\}$ of transition-consistent and volume-normalized measures on $\mathcal{P}_n$. We abuse notation and consider the set of n-supertiles to be the interval $[(3/2)^n, 3(3/2)^n]$, in which case the measures of sets of supertiles is given by $\rho_n([a,b]) = \int_a^b f_n(x)dx$, where dx is ordinary Lebesgue measure. The only choice allowing a transition-consistent and volume-normalized sequence is

$$f_n(x) = \begin{cases} \dfrac{1}{(3\ln 3 - 2\ln 2)x^2} & (3/2)^n \le x \le 2(3/2)^n \\ \dfrac{3}{(3\ln 3 - 2\ln 2)x^2} & 2(3/2)^n < x \le 3(3/2)^n. \end{cases}$$

Although it is not obvious, it is true that the system is uniquely ergodic.

The invariant measure is invariant under scaling in the sense that $\mu(\Omega_{I,U}) = \mu(\Omega_{\lambda I,\lambda U})$ for a trim set I and denoting by λI the set of patches obtained by rescaling and subdividing I. This probably implies that the measure is absolutely continuous.

6.1.4. Transverse topology. As usual for tilings of the line, the transversal $\Xi(\Omega)$ is the set of all tilings that have a tile endpoint at the origin. Although it appears that the transversal might be connected, it is in fact a Cantor set.

To show that any given tiling $\mathbf{T} \in \Xi(\Omega)$ is not isolated, consider any $\epsilon > 0$ and an N-supertile in $\mathbf{T}$ that contains $B_{2/\epsilon}(0)$. (Even though this supertile may not be a natural supertile of $\mathbf{T}$, it must exist since $\mathbf{T}$ is admitted by the substitution rule). There are lots of supertiles that are within $\epsilon/2$ of this supertile in the patch metric, since we can choose a supertile as close to the original in length as necessary to ensure they subdivide to comparable tiles. Select one and let $\mathbf{T}' \in \Xi(\Omega)$ be a tiling with this supertile at the origin. Then $d(\mathbf{T}, \mathbf{T}') < \epsilon$.

To show that the transversal is totally disconnected consider two tilings $\mathbf{T}, \mathbf{T}' \in \Omega$ that are close, so that they very nearly have the same patch of tiles in a large ball B around the origin. All the tiles in $\mathbf{T}$ are multiples $2^j/3^k$ of each other, as are those in $\mathbf{T}'$. Thus in B all the tiles in each are multiples by the *same* $2^j/3^k$ of their respective tiles at the origin. Suppose the tile for $\mathbf{T}$ at the origin is slightly larger than that in $\mathbf{T}'$. All the tiles in B in $\mathbf{T}$ are then larger than those of $\mathbf{T}'$. Outside of B there will be a tile in $\mathbf{T}'$ that is of size $3 - \epsilon$ such that in the corresponding region in $\mathbf{T}$ the 1-supertile is of size slightly larger than 3 and thus is broken into two tiles of sizes slightly greater than 1 and 2. Using this difference we can partition the transversal into two clopen sets: one containing all tilings with a tile of size less than or equal to 3/2 at the left side of this location and one containing all tilings with a tile of size greater than 3/2 there.

6.1.5. Complexity. Let $N(\epsilon, L)$ be the complexity function that counts the minimum number of patches of d_L-radius ϵ it takes to cover Ω. Recall that a d_L-ball of radius ϵ is the set of all tilings that have a patch that agrees with a fixed tiling on

$[-1/\epsilon, L+1/\epsilon]$ up to ϵ. We concern ourselves with the situation where L is much greater than $1/\epsilon$.

In order to specify the fixed tilings we use as the 'centers' of the balls, we need to specify up to ϵ all patches of size $L+2/\epsilon$. To do so, we need to find the smallest supertile that contains an interval of that size. The length of that supertile is on the order of L and we need to specify it up to ϵ, so there are on the order of L/ϵ choices for that. We also need to specify precisely where within the supertile we are up to ϵ, and that gives us on the order of L/ϵ choices also. This means that the complexity goes as L^2/ϵ^2. This is polynomial complexity, and the ϵ-entropy is zero.

6.2. Example 6: Direct product variation

We begin our analysis of the DPV tiling space Ω with an alternative description of how to obtain it by projection from a structure in $\mathbb{R}^3$. Then we show that it is minimal and show how to compute the invariant measure. The topological spectrum and complexity are computed and the transversal is discussed for varying values of tile widths. We conclude the section with a brief description of how the discussion would extend to other DPVs.

6.2.1. Projection method. A standard trick with one-dimensional substitution sequences is to construct the so-called 'broken line' or 'staircase' of a substitutive sequence. The staircase lives in $\mathbb{R}^{|\mathcal{A}|}$ and is constructed iteratively by starting at the origin, taking a step in the $\vec{e}_{x(0)}$ direction, then one in the $\vec{e}_{x(1)}$ direction, and so on. It is well known that such a staircase will approximate the Perron eigenline of the substitution matrix, and projection of the staircase onto the Perron eigenline along the weak eigendirections yields the natural tile lengths for the corresponding self-similar tilings. Projection of the staircase onto other lines or in other directions produces different tilings that may or may not be conjugate to each other, depending mostly (it seems) on the expansion factor of the system. When the expansion matrix is Pisot all the points on the staircase lie within a bounded distance of the eigenline. When it is not the staircase can wander arbitrarily far away from the eigenline, but it comes close to the eigenline repeatedly.

In a planar tiling context there are similar constructions when the expansion constant is Pisot [1] or if other special conditions hold [2]. In those cases one constructs a 'discrete plane' made up of two-dimensional facets in some $\mathbb{R}^n$. The conditions in our examples do not satisfy these conditions, but nevertheless it is possible to construct a discrete surface that projects onto our DPV tilings. This surface has holes that cannot be seen when we project in a direction that forms a tiling. We explain the method for our basic example here and a description of how to generalize it appears at the end of this section.

The stepped surface will appear in $\mathbb{R}^3$ and can be constructed by substitution. The $A = a \times c$ tile corresponds to the unit square spanned by the origin, $(1,0,0)$ and $(0,0,1)$. The $B = b \times c$ tile corresponds to the unit square spanned by the origin, $(0,1,0)$, and $(0,0,1)$. The substitution rule is pictured in Figure 3.

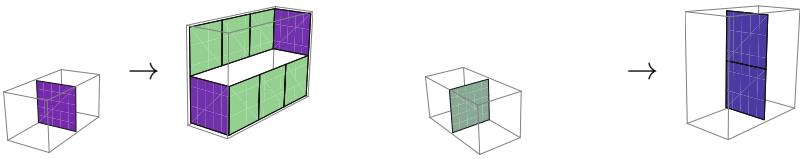

FIGURE 3. The facet substitution.

It can be proved that the substitution can be iterated indefinitely as long as one follows concatenation rules; alternatively the n-supertiles can be seen as fusions of the $(n-1)$-supertiles according to the same basic combinatorics as the original fusion. In Figure 4 we show a 3-supertile at an angle different than one of the projection angles we will use to see the planar tiling. From this angle the holes in the surface are visible, and since the facet of the hypercube corresponding to an $a \times b$ is not an element of our tile set the holes are parallel to the xy-plane. To obtain our planar DPV with parameters a, b, c, we simply take the stepped surface and project it using the matrix $\begin{pmatrix} a & b & 0 \\ 0 & 0 & c \end{pmatrix}$.

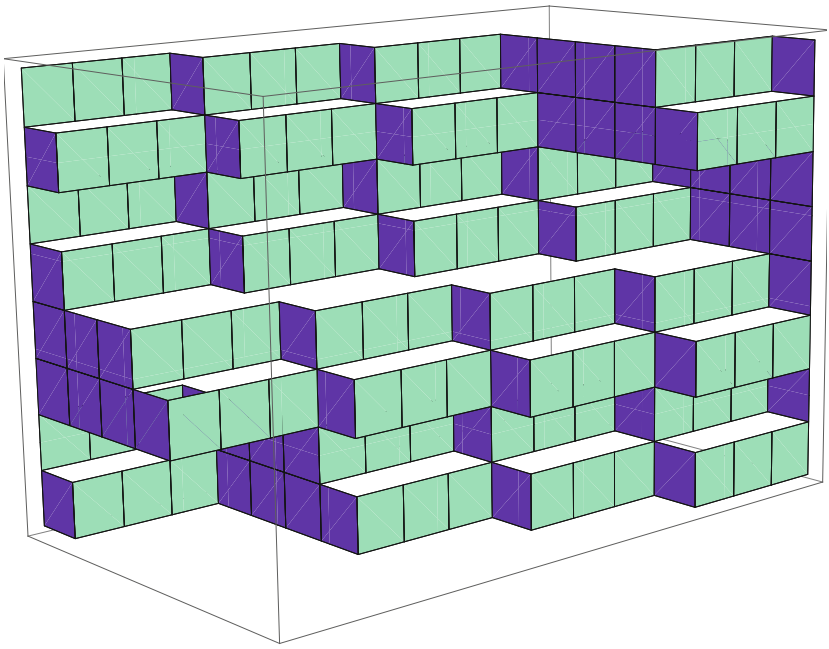

FIGURE 4. The 3-supertile of type $a \times c$.

Now we begin to see why the fault lines occur. The holes correspond to pieces of fault line and we see a large hole halfway up the supertile that will continue to grow as we iterate the substitution. In the limit the top and bottom halves become

totally severed and this is reflected in Ω by tilings with arbitrary shifts along the fault line.

6.2.2. Minimality and invariant measure. Our example has transition matrix $M = \begin{pmatrix} 2 & 2 \\ 6 & 0 \end{pmatrix}$ as a standard substitution rule and so $A_{n,N} = M^{N-n}$. Since M is a primitive matrix we immediately see that the fusion rule must be primitive and so its dynamical system is minimal.

Using results from [12] we know that for any width parameters a and b the dynamical system will be uniquely ergodic. The frequency measure on patches is atomic and gives all patches that live in some n-supertile a positive frequency. All patches that are admitted only in the limit and never appear in any n-supertile have measure 0.

The precise nature of the ergodic measure depends on a and b for several reasons. First off, if a and b are rationally related then the tiling space has finite local complexity. The least common denominator q determines the horizontal offsets, and all all multiples of $1/q$ occur. This means that the patches that we are measuring the frequency of are different as we change the parameters. If a and b are rationally independent then we have infinite local complexity occurring along horizontal fault lines (i.e., there is a fracture in the direction of $(1, 0)$). Patches with frequency 0 appear.

The Perron eigenvalue is $\lambda = 1 + \sqrt{13}$, which is the product of the length expansions of the one-dimensional substitutions, and it has left eigenvector $\vec{v} = (\lambda, 6)$. This vector, when normalized by volumes of $P_n(A)$ and $P_n(B)$, gives the frequencies $\rho_n(A)$ and $\rho_n(B)$ of those supertiles. The volume of $P_n(A)$ or $P_n(B)$ can be computed in two ways for a general a and b. Since $\mathrm{Vol}(A) = ac$ and $\mathrm{Vol}(B) = bc$, we see that $(\mathrm{Vol}(P_n(A)), \mathrm{Vol}(P_n(B))) = (ac, bc) * M^n$. On the other hand the volume of $P_n(A)$ or $P_n(B)$ is $2^n c$ times the length of the nth one-dimensional substitution of the a or b tile. To find ρ_n, we see that it is $(\lambda, 6)/K$, where $K = \lambda P_n(A) + 6P_n(B)$. In the special case where a and b are natural tile lengths for the horizontal substitution, the tiling is self-affine and K is a constant multiple of λ^{-n}.

6.2.3. Topological Spectrum. No matter what our choice of a and b are there will be eigenvalues of the form $(0, c\mathbb{Z}[1/2])$. This is because there are functions that detect only where a tiling is in the vertical, solenoid hierarchy. Whether or not there is any horizontal spectrum depends on whether a and b are rationally related.

If a and b are irrationally related, the fault lines provide us with fractures of the form $(k, 0)$ for arbitrarily small k. By Corollary 5.6 we see then that any eigenvalue must be of the form $(0, y)$, so the spectrum we retain is that of the solenoid.

If $a = pb/q$, then we obtain additional discrete spectrum of the form $(kp/a, 0)$, for $k \in \mathbb{Z}$. The eigenfunctions keep track of the horizontal offsets of, say, the tile containing the origin. $f(\mathbf{T})$ would equal the displacement of $\mathbf{T}$ horizonally from the nearest $1/q$ piece of the tiling.

6.2.4. Transverse topology. The topology of the transversal depends on whether a and b are rationally related. If they are, then the tiling space has FLC and the transversal is a Cantor set. However, when a and b are not rationally related the topology of the transversal is not well understood. We show that the transversal is not a Cantor set and that in particular it is not totally disconnected. For concreteness, assume the control points of the tiles are their centers of mass.

Any tiling in the transversal that has a fault line is in a connected component of the transversal that contains all possible shifts of the half-plane that does not contain the origin. Consider a point $\mathbf{T} \in \Xi(\Omega)$ that does not have a fault line. It is arbitrarily close to tilings that do have fault lines: any ball around the origin is contained in some large supertile of $\mathbf{T}$, and that supertile can be contained in a tiling with a fault line. Thus $\mathbf{T}$ exactly agrees with a continuum of tilings in $\Xi(\Omega)$ on this supertile.

In this example there are infinitely many connected components that are distinguished from one another by where the tiling is in the vertical, solenoid hierarchy. However it is not clear whether a location in the hierarchy contains multiple connected components that are distinguished in some way by the horizontal substitution. Understanding the transversal for tilings with both vertical and horizontal fault lines is even more mysterious.

6.2.5. Complexity. The complexity of Ω will depend on the parameters a and b. In the case where a and b are rationally related and Ω has FLC we get a tiling space that has higher complexity than a planar self-similar tiling because of the unbounded nature of the holes. For simplicity suppose a and b are integer multiples of $1/q$, for $q \in \mathbb{N}$. Suppose that L is greater than $1/\epsilon$ and both are greater than q. We need to count how many patches of size L there are, up to ϵ. There is some minimum size of supertile for which every patch of size L is contained entirely within a supertile or across the boundary of 2, 3, or 4 supertiles. There are only a finite number of choices for such patterns since we only have two supertile types. Inside of any such pattern we have a bounded multiple of L^2/ϵ^2 places to put the corner of a patch; moreover for any patch that has two or more supertiles there are about qL ways they can fit together. This implies that the complexity as $L \to \infty$ is bounded above and below by bounded multiples of qL^3/ϵ^2.

Now suppose instead that a and b are irrationally related. Most of the argument is the same except that now there are L/ϵ ways to fit two or more supertiles of size L together. This leaves us with a complexity bounded above and below by bounded multiples of L^3/ϵ^3. In either case we have polynomial complexity $C(\epsilon)L^3$, with the exponent 3 being notable because it is greater than the ambient dimension of the tiling spaces. Although the system still has entropy 0, it is more complex than planar FLC self-similar tilings.

6.2.6. Other DPVs. Any DPV can be seen as the projection of a higher-dimensional stepped surface, but we will keep our discussion in the plane for simplicity. If the alphabets are $\mathcal{A}$ and $\mathcal{B}$, then there are $|\mathcal{A}||\mathcal{B}|$ tile types, all of the form

$a_i \times b_j$. The unit hypercube in $\mathbb{R}^{|\mathcal{A}|+|\mathcal{B}|} = \mathbb{R}^{|\mathcal{A}|} \times \mathbb{R}^{|\mathcal{B}|}$ has (($|\mathcal{A}| + |\mathcal{B}|$) choose 2) two-dimensional facets that contain the origin. We need $|\mathcal{A}||\mathcal{B}|$ of them for our stepped surface, which we construct as follows. Let the standard basis vectors in $\mathbb{R}^{|\mathcal{A}|}$ correspond to the letters in $\mathcal{A}$ and let the standard basis vectors of $\mathbb{R}^{|\mathcal{B}|}$ correspond to those in $\mathcal{B}$, so that the tile $a_i \times b_j$ is represented by the corresponding facet that contains the origin, one standard direction in $\mathbb{R}^{|\mathcal{A}|}$, and one in $\mathbb{R}^{|\mathcal{B}|}$. The remaining hypercube faces that contain the origin are those that take both other vectors in either $\mathbb{R}^{|\mathcal{A}|}$ or $\mathbb{R}^{|\mathcal{B}|}$. These don't correspond to tiles in the DPV and so do not appear in the facet substitution (or rather, they appear as holes). To give the tiling parameters $a_1, a_2, \ldots, a_{|\mathcal{A}|}, b_1, \ldots, b_{|\mathcal{B}|}$ we simply project it to the plane via the matrix $\begin{pmatrix} a_1 & \cdots & a_{|\mathcal{A}|} & 0 & \cdots & 0 \\ 0 & \cdots & 0 & b_1 & \cdots & b_{|\mathcal{B}|} \end{pmatrix}$.

It is easy to construct DPVs that have fractures in both the horizontal and vertical directions by varying the direct product of two substitutions that have non-Pisot length expansions. (The primary example in [9] is of this form.) By choosing irrationally related side lengths and irrationally related height lengths we obtain a topologically weakly mixing system.

6.3. Example 7: Solenoid extensions

Here we consider the tiling space Ω_c obtained by the compact label set $\mathcal{L} = \mathbb{N}_0 \cup \mathcal{L}_c$. For convenience of notation we will make a partial order on the label set: $k < m$ whenever that is true for k and m as integers, and also if k is an integer and $m \in \mathcal{L}_c$. Every integer label is 'less than' every compactification label. We will see that on the measurable level the Ω_cs are all the same, but on the topological level they are quite different.

6.3.1. Minimality. By looking at $P_4(k)$ for some $k \geq 4$ we can think about the transition map $A_{n,N} : \mathcal{P}_n \to \mathcal{P}_N$.

$$P_4(k) = P_3(k)P_3(3) = k0102010\,30102010.$$

We have one 3 and one k; a pair of 2s, four ones, and eight zeroes. If instead we break it into 1-supertiles, we have four of type 1, two of type 2, and one each of types 3 and k. We begin to suspect that the number of n-supertiles of type m, for $m \geq n$, in an N-supertile is equal to the number of tiles of type m and has nothing to do with N as long as N is larger than m. If m is larger than N, then it is 0 unless we are in an N-supertile that is also of type m. In the former case we see that the number of n-supertiles of type m in $P_N(k)$ is $2^{N-n}/2^{m-n+1}$ and we have

$$A_{n,N}(m,k) = \begin{cases} 2^{N-(m+1)} & \text{if } n \leq m < k \text{ and } m < N \\ 1 & \text{if } N \leq m = k \\ 0 & \text{otherwise.} \end{cases}$$

For a fixed n and $\epsilon > 0$, we can always choose an N such that every N-supertile contains a tile arbitrarily close to $P_n(m)$ for any $m \geq n$. The details of this

will depend on the topology of $\mathcal{L}$, in particular on the speed with which sequences converge to elements of $\mathcal{L}_c$. This proves primitivity and hence minimality.

6.3.2. Measurable conjugacy of Ω_c to the dyadic solenoid. Every solenoid extension is measurably conjugate to the dyadic solenoid, defined as the inverse limit $\mathbb{S} = \varprojlim(S^1, \times 2)$. Elements of $\mathbb{S}$ take the form $(x_0, x_1, x_2, \dots)$ such that for each n, $2x_n \equiv x_{n-1} \mod 1$. We define $f(\mathbf{T}) = (x_0, x_1, x_2, \dots)$ as follows. The position of the origin in the tile containing it determines x_0, choosing $x_0 = 0$ if the origin is on the boundary between two tiles. The origin is either in the left or right half of the 1-supertile that contains it. If it is in the left half we let $x_1 = x_0/2$ and if it is in the right half or exactly in the middle we let $x_1 = (x_0 + 1)/2$. Once this is determined we look at whether the origin is in the left half of the 2-supertile containing it or if it is in the right (or exactly in the middle). In the former case we let $x_2 = x_1/2$ and in the latter we let $x_2 = (x_1 + 1)/2$. Since the origin lies somewhere in an n-supertile for every $n \in \mathbb{N}_0$ we can inductively define f for all of Ω_c.

Every tiling $\mathbf{T} \in \Omega_c$ that doesn't contain any element of $\mathcal{L}_c$ is sent to a unique point in $\mathbb{S}$. Since such tilings form a set of full measure we have measurable conjugacy. The rest of the tilings in Ω_c do not map in a one-to-one fashion since a tiling that is made of two infinite-order supertiles maps to an element of the solenoid that knows the position of the origin in its supertile but not what element of $\mathcal{L}_c$ the tiling contains. This also explains why the one-point compactification is topologically conjugate to the solenoid.

Since Ω_c is measurably conjugate to the dyadic solenoid it must be uniquely ergodic and share its purely discrete measurable spectrum $\mathbb{Z}[1/2]$. The frequency measures are supported on the subset of the transversal on which f is one-to-one, i.e., no infinite-order supertiles. However, the nature of the limits does affect the nature of the measurable isomorphism. A ball of radius ϵ around a tiling with a given element of $\mathcal{L}_c$ at the origin maps onto a set of equal measure in the solenoid, but this set will vary as the elements of $\mathcal{L}_c$ do, or when the sequences that converge to them do.

6.3.3. The topological impact of $\mathcal{L}_c$. To begin we consider the one-point compactification that has elements of $\mathbb{N}_0$ converging monotonically to a single limit point ℓ. In this case we can see that Ω_c is topologically conjugate to the dyadic solenoid using the argument from the previous section on measurable conjugacy. In that argument the map is one-to-one for each tiling in Ω_c that does not contain an element of $\mathcal{L}_c$ but multiple-to-one on those that do. However, the multiple copies are in correspondence to the number of elements of $\mathcal{L}_c$. Since in this case there is only one element of $\mathcal{L}_c$ the map continues to be one-to-one on the exceptional tilings, so it is one-to-one everywhere and thus provides a topological conjugacy to the dyadic solenoid.

This gives us the unusual situation where the tiling dynamical system is not expansive. To see this, choose any $\delta > 0$ and find an N such that if $N', N'' > N$, then the distance between prototiles labelled N' and N'' is less than δ. Any two

tilings T and T' for which the origin sits in the same location in its N-supertile will then satisfy the property that $d(T - x, T' - x) < \delta$ for all $x \in \mathbb{R}$ regardless of the type of the N-supertile. The way to think about this is that each N-supertile differs from any other only on the first tile, and those first tiles are labelled greater than N and thus differ by less than δ.

The situation for two-point compactifications becomes more subtle. The version for which all even numbers converge monotonically to ℓ and all odd numbers converge monotonically to ℓ' is shown to be topologically conjugate to the FLC tiling space generated by the 'period-doubling' substitution $X \to YX, Y \to XX$, where X and Y are unit length tiles. If instead $\mathbb{N}_0$ is partitioned into two sets S and S', one converging to ℓ and the other to ℓ', the measures of clopen subsets of the transversal are determined by $\alpha = \sum_{n \in S} 2^{-n}$, and thus so is the gap-labelling group. Two tiling spaces, one with α and the other with α' can be compared, and whether or not they are topologically conjugate or even homeomorphic depends on how these constants are related. See [12] for details.

By compactifying with limit sets that have an interesting topology, we can obtain tiling spaces with nontrivial cohomologies in H^n. These examples can be distinguished by their cohomology.

6.3.4. Transverse topology. The topology of the transversal depends on the topology of $\mathcal{L}_c$. Suppose that $\mathbf{T}, \mathbf{T}' \in \Xi(\Omega)$ are two tilings that have an element of $\mathcal{L}_c$ at the origin. Note that $\mathbf{T}$ and $\mathbf{T}'$ are made of two infinite-order supertiles and thus can only differ at the origin. So if it were possible to make two open sets in $\Xi(\Omega)$ that disconnected the transversal and contain one each of $\mathbf{T}$ and $\mathbf{T}'$, those sets would correspond to two open sets in $\mathcal{L}_c$ that disconnected $\mathcal{L}_c$. Conversely, if $\mathcal{L}_c$ can be separated by two open sets, then if $\mathbf{T}$ and $\mathbf{T}'$ contain an element of these sets at the origin, we can partition the rest of the elements of $\mathcal{L}$ so that they are close to one or the other of the subsets of $\mathcal{L}_c$ to create a pair of open sets that disconnect $\Xi(\Omega)$ and contain $\mathbf{T}$ and $\mathbf{T}'$ respectively. Thus $\Xi(\Omega)$ is disconnected if and only if $\mathcal{L}_c$ is.

In particular, the one- and two-point compactifications of the solenoid have totally disconnected transversals. The one-point compactification is not strongly expansive because it is not expansive. To see this, consider a fixed $\delta > 0$ and let $d(\mathbf{T}_1, \mathbf{T}_2) \leq \delta$, where $\mathbf{T}_1$ and $\mathbf{T}_2$ have an N-supertile beginning at the origin, where N is sufficiently large that any two elements of $\mathcal{L}_c$ that are greater than N are within δ in the tile metric. Then $\mathbf{T}$ and $\mathbf{T}'$ differ at most on the beginning of each N-supertile, but these are all within δ of one another in the patch metric. Thus $d(\mathbf{T} - k, \mathbf{T}' - k) < \delta$ for all $k \in \mathbb{Z}$ and the dynamics on the transversal are not expansive (and indeed the overall dynamical system is not expansive). Thus Theorem 5.9 confirms that the one-point compactification is not homeomorphic to any FLC tiling system.

By a similar argument we can show that the two-point compactification is homeomorphic to a FLC tiling space by Theorem 5.9. This is because even if two

tilings very nearly agree on the central N-supertile, each of them are bound to have N-supertiles whose first tiles are close to distinct elements of $\mathcal{L}_c$ and thus not to each other.

If $\mathcal{L}_c$ is not totally disconnected, then its solenoid extension is not homeomorphic to a one-dimensional tiling space with finite local complexity.

6.3.5. Complexity. This is the topological invariant that lets us see that the one-point compactification, while failing to have finite local complexity, does so in a way that is much less complex than a non-periodic FLC tiling space. In order to be specific about complexity let us use $N_1(\epsilon, L)$, the minimum number of d_L- balls of radius ϵ balls it takes to cover Ω. Take $N \in \mathbb{N}_0$ such that any element $n \geq N$ is within ϵ of the limit point ∞ in $\mathcal{L}$. Every n-supertile with $n \geq N$ is within ϵ of any other in the patch metric. This means that the system is basically periodic with period 2^N, up to ϵ, and $N_1(\epsilon, L)$ is bounded multiple of 2^N once L is sufficiently large relative to ϵ. This means that the system has bounded complexity in the sense of Section 3.4.

No matter how complicated $\mathcal{L}_c$ is, the complexity cannot become too high and there will always be zero entropy. This is because, for any $\epsilon > 0$, an ϵ-cover of $\mathcal{L}$ into N elements effectively reduces Ω_c to a solenoid extension on N elements.

7. Important questions

There are two main types of questions that seem to be important in the study of tiling spaces with infinite local complexity. One is to ask, what properties of these spaces are invariant under homeomorphism or some other type of conjugacy? Another is to ask, how do combinatorial and geometric factors influence the dynamical, measure-theoretic, topological, or complexity properties?

It was already known from [20] that infinite local complexity is not a topological invariant. The solenoid extensions considered here show that infinite local complexity is not preserved by measurable or topological conjugacy either. So what are the important classes of infinite local complexity? Even though FLC is not an invariant property, there should be some properties that guarantee that a tiling space is topologically conjugate, or measurably conjugate, or just homeomorphic to one that is FLC. In one dimension, Theorem 5.9 gives a topologically-invariant characterization of this class. In higher dimensions geometry becomes an obstacle and a direct generalization is unlikely to be sufficient: the methods used in the one-dimensional proof don't produce valid tilings in higher dimensions. So the topological classification of FLC in 2 or higher dimensions remains unresolved, along with the dynamical classification in any dimension.

The property of being constructed from a finite set of prototiles is also not invariant. A simple example is the tiling space made out of unit square tiles that lie in rows that are offset by random amounts. We can introduce an infinite number of tile types symbolically, by constructing an infinite label set given by all the possible coronas, which produces a topologically conjugate tiling space. Alternatively, we

can introduce an infinite number of tile types geometrically by taking the space of dual tilings. If we use the centers of the tiles as the vertices of the dual prototiles, the dual prototile set consists of one square and infinitely many triangles of equal areas but different angles, and is not even compact. Since many important tiling models assume a finite prototile set, knowing properties that guarantee that a space is equivalent to one with this property seems essential.

The combinatorics in the tilings in the last paragraph are all basically the same and depend only on whether an offset between rows is zero or not. Thus the combinatorics of that particular tiling space do not reveal information about local complexity. By way of contrast, the combinatorics of the DPV of Example 6 will determine whether the tiling space has finite or infinite local complexity. The interplay between combinatorics and complexity is especially evident in the DPV case. If one does not vary the direct product substitutions, no geometric factors can influence the local complexity of the tiling space: it is guaranteed to have FLC. But if one does vary the direct product to produce a DPV, then the local complexity depends on several factors: the inflation constant, the specific combinatorics of the substitution, and the sizes of the tiles.

More generally, in the category of primitive fusion or substitution tilings the interplay between combinatorics, geometry, number theory, topology, and dynamics is an important area of investigation. It is known that some form of Pisot condition on the expansion factor has a profound effect on dynamics and complexity (cf. [22, 9, 5]). Such Pisot conditions enforce a rigidity on the tiling space; the dynamics and complexity are severely restricted. In the absence of a Pisot condition, combinatorics and geometry can influence the tiling space in numerous ways. In order to understand certain tiling models it is important that we understand the nature of this influence.

References

[1] P. Arnoux, V. Berthé, and A. Siegel, Two-dimensional iterated morphisms and discrete planes, *Theor. Comp. Sci.* **319**, (1-3), (2004), 145–176.

[2] P. Arnoux, M. Furukado, E. Harriss, and S. Itô, Algebraic numbers, free group automorphisms and substitutions on the plane. *Trans. Amer. Math. Soc.* **363**, no. 9 (2011), 4651–4699.

[3] M. Baake and U. Grimm, Aperiodic Order, Volume 1: A Mathematical Invitation, in *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge, 2013.

[4] M. Barge and C. Olimb, Asymptotic structure in substitution tiling spaces, *Ergod. Th. Dyn. Sys.* **FirstView**, DOI 10.1017/etds.2012.118, http://journals.cambridge.org/article_S0143385712001186, (2013), 1–40.

[5] A. Clark and L. Sadun, When size matters: subshifts and their related tiling spaces, *Dyn. Syst.* **23** (2001), 1043–1057.

[6] L. Danzer, Inflation species of planar tilings which are not of locally finite complexity. *Proc. Steklov Inst. Math.* **230** (2002), 118–126.

[7] T. Downarowicz, Survey of odometers and Toeplitz flows, *Contemporary Mathematics* **385** (2005) 7–37.

[8] N.P. Frank, A primer on substitutions tilings of Euclidean space, *Expo. Math.* **26**, no. 4 (2008), 295–326.

[9] N.P. Frank and E.A. Robinson, Jr., Generalized β-expansions, substitution tilings, and local finiteness, *Trans. Amer. Math. Soc.* **360**, no. 3 (2008), 1163–1177.

[10] N.P. Frank and L. Sadun, Topology of some tiling spaces without finite local complexity, *Contin. Dyn. Syst.* **23**, no. 3 (2009), 847–865.

[11] N.P. Frank and L. Sadun, Fusion: a general framework for hierarchical tilings of $\mathbb{R}^d$, *Geom. Dedicata*, DOI 10.1007/s10711-013-9893-7, http://dx.doi.org/10.1007/s10711-013-9893-7, (2013) 1–38.

[12] N.P. Frank and L. Sadun, Fusion tilings with infinite local complexity, *Topology Proc.* **43** (2014), 235–276.

[13] E. Harriss and D. Frettlöh, *Tilings encyclopedia*, http:/tilings.math.uni-bielefeld.de/

[14] C. Goodman-Strauss, Matching rules and substitution tilings, *Annals of Math.* **147** (1998), 181–223.

[15] A. Julien and J. Savinien, Transverse Laplacians for substitution tilings, *Comm. Math. Phys.* **301**, no. 2 (2011), 285–318.

[16] R. Kenyon, Rigidity of planar tilings, *Invent. math.* **107** (1992), 637–651.

[17] N. Pytheas Fogg, Substitutions in Dynamics, Arithmetics, and Combinatorics, *Lecture Notes in Mathematics* **1794**, Springer-Verlag, Heidelberg, 2002.

[18] C. Radin, The pinwheel tilings of the plane, *Annals of Math.* **139:3** (1994), 661–702.

[19] C. Radin, *Miles of Tiles*, Student Mathematical Library **1**, American Mathematical Society, Providence (1999).

[20] C. Radin and L. Sadun, Isomorphism of hierarchical structures, *Ergod. Th. Dyn. Sys.* **21** (2001), 1239–1248.

[21] L. Sadun, Some generalizations of the pinwheel tiling, *Disc. Comp. Geom.* **20** (1998), 79–110.

[22] B. Solomyak, Dynamics of self-similar tilings, *Ergod. Th. Dynam. Sys.* **17** (1997), 695–738.

Natalie Priebe Frank
Department of Mathematics
Vassar College
Poughkeepsie, NY 12604, USA
e-mail: nafrank@vassar.edu

Progress in Mathematics, Vol. 309, 259–306
© 2015 Springer Basel

On the Noncommutative Geometry of Tilings

Antoine Julien, Johannes Kellendonk and Jean Savinien

Abstract. We discuss constructions of spectral triples for spaces which are related to subshifts or aperiodic tilings. These spectral triples give rise to distance functions, zeta functions, Laplace operators and K-homology classes. We investigate these objects, and relate their properties to standard notions in tiling theory.

Mathematics Subject Classification (2010). 46L87, 37B80, 54H20.

Keywords. Spectral triples, subshifts, tilings.

1. Introduction

Alain Connes' noncommutative geometry program is based on translating the ordinary notions of geometry into the framework of associative algebras in such a way that they make sense also for noncommutative algebras. This is very natural from the point of view of quantum physics. The basic objects are then no longer spaces but algebras.

Methods from noncommutative topology – a branch of noncommutative geometry – have been applied quite early to aperiodic tilings. Although noncommutative topology will be touched on only briefly in the last section of this article we will shortly describe the history of its applications to tilings of the Euclidean space. The start was made by Alain Connes himself and by Jean Bellissard. In fact, one of the first examples of a "noncommutative space" which Connes discussed in his book [11] is a description of the most famous of all aperiodic tilings in the plane, the Penrose tilings, and Jean Bellissard proposed a C^*-algebraic approach to the description of aperiodic media [2].

The point of departure in a noncommutative theory is an associative algebra, often a C^*-algebra, which, in the context we consider here, should be somehow derived from a tiling or a point set. Connes' version of the C^*-algebra for the Penrose tilings is based on their substitution rule. Robinson had already observed that the substitution rule can be used to describe the set of Penrose tilings modulo isometry as the set of $\{0, 1\}$-sequences without consecutive 1's [25], a description

which Connes recognized as the set of paths on the Bratteli diagram which has inclusion graph A_4. Consequently, Connes assigned to the Penrose tiling the AF-algebra which is defined by the Bratteli diagram.

Bellissard's C^*-algebra of an aperiodic medium is the crossed product algebra defined by a covariant family of Schrödinger operators associated with given local atomic potentials, a technique which had been devised for the study of disordered systems. Such families are representations of the crossed product algebra of a dynamical system whose underlying space is what Bellissard introduced as the *hull of the potential* and whose group action is given by translation of the potential in space (see [5] for a later review on the subject). The spatial arrangement of a medium may be described by a tiling or a point set and thus Bellissard's crossed product algebra may be considered as a version of a C^*-algebra for the tiling.

Later, one of the authors provided a direct geometric construction of an algebra for a tiling, the so-called discrete tiling algebra. The algebra is the groupoid-C^*-algebra of an étale groupoid which arises when considering the translations which one can make in order to go from the center of one tile of the tiling to the center of another tile [30]. No Schrödinger operator is needed for this construction, neither a substitution rule. In its most rigorous form in can be derived from a purely combinatorial object, the inverse semigroup of the tiling, the topology being entirely determined by the order relation on the semigroup [31].

Finally, Ian Putnam with his student Jared Anderson constructed an algebra for a tiling again in form of the crossed product algebra of a dynamical system. This dynamical system had been originally defined by Rudolph [55]. Its underlying space is the *continuous hull of the tiling* and the algebra is referred to as the continuous tiling algebra. The continuous hull is closely related to the hull of the potential and often the same.

To summarise, there were several approaches to construct the basic object for the noncommutative geometry of tilings. The first was derived from a substitution, the second from a potential whose spatial repetition is governed by the tiling, the third and the fourth directly from the tiling. As it turned out, the second and the fourth algebra are essentially equal and can be seen as the stabilized version of the third algebra. From the point of view of noncommutative topology this means that the three latter approaches are equivalent. Connes' AF-algebra is closest to the third, the discrete tiling algebra, namely it is a (proper) subalgebra of the latter. The noncommutative topological invariants of the tiling algebra are richer than those of this AF-algebra. The AF-algebra should be regarded as describing the substitution symmetry of the tiling rather than the tiling itself.

Newer developments concern the noncommutative topology of tilings with infinite rotational symmetry, like the Pinwheel tilings [45], of tilings in hyperbolic space [47] and of combinatorial tilings [51, 52]. These developments will not be explained in this book.

All the above algebras are noncommutative C^*-algebras and come with interesting maximal commutative subalgebras. These are the algebra of continuous

functions over the continuous hull, the algebra of continuous functions over an abstract transversal of that hull (often referred to as canonical transversal or discrete hull), and the commutative subalgebra of the AF-algebra which is the fixed point algebra of the latter under rotations and reflections.

The next step in noncommutative topology is the investigation of the K-theory of the algebras. This has been to a large extend discussed in the review article [34]. While it is quite simple to compute the K-groups for AF-algebras, it is less so for crossed products and the computation of the K-theory for the discrete or continuous tiling algebra (both yield the same answer) forms a substantial part of tiling theory. In fact, neglecting order the K-groups are isomorphic to cohomology groups (at least in low dimension with integral coefficients, and always with rational coefficients), and the latter are computable for certain classes of tilings, namely substitution tilings and almost canonical cut-and-project tilings. This is described in Chapters 3 and 4 of this book.

Another part of noncommutative geometry is cyclic cohomology and/or K-homology and their pairing with K-theory. Here the older results cover only the pairing with very specific cyclic cocycles, like the trace, the noncommutative winding number and the noncommutative Chern character. We will have to say something more systematic about the pairing of K-homology with K-theory in the last section of this chapter.

We end this short survey on the noncommutative topology of tilings with the remark that both, noncommutative geometry and the theory of tilings were strongly motivated and influenced by physics. The K-groups and their pairing with cyclic cocycles have relevance to physics, namely to topological quantization, in particular in the gap-labelling [5, 2, 3], the Integer Quantum Hall Effect [4, 38] and other topological insulators [46], the pressure on the boundary [32] and the Levinson's theorem [35, 36, 37, 8]. These fascinating developments are also beyond the scope of this book.

After this short description of the developments of the noncommutative topology of tilings we come to the topic of this article, namely the noncommutative *geometry* of tilings. Geometry is the investigation of a space through the measurement of length or distance between points. The fundamental object in geometry is thus a length or distance function. In noncommutative geometry it is hence a notion which allows one to construct the analog of such a function for (possibly noncommutative) algebras, the guiding principle being always the duality between C^*-algebras and topological spaces. Motivated by quantum physics Connes advocates that length should be a spectral quantity (we won't try to be precise on what that really means). He proposed the notion of a spectral triple as the fundamental object of noncommutative geometry.

Leaving the precise definition for later (Definition 2.1) we can say that a spectral triple is a representation of our favorite associative algebra on a Hilbert space together with a choice of a self-adjoint typically unbounded operator on that Hilbert space, the Dirac operator D, such that $[D, \cdot]$ can play the role of a

derivative. The notion is modeled after Riemannian spin manifolds and then cast into an axiomatic framework which we will partly recall further down. The subject of this chapter is to summarize what is known about spectral triples for tilings. This subject is wide open, in fact, no satisfying spectral triple has been constructed for the full tiling algebra (neither the discrete, nor the continuous one). What has been achieved so far are constructions for the maximal commutative subalgebras mentioned above. Though this seems disappointing at first sight, it has already led to new concepts for the study of tilings.

What is a spectral triple good for? As we said already, it provides us with a noncommutative notion of distance, namely a pseudometric on the space of states of the algebra. In the case we consider here, a commutative tiling algebra, this yields in particular a distance function on the discrete or the continuous hull. This is an extra structure and we may ask: Does this distance function generate the topology? And if not, what does this say about the tiling? A second ingredient a spectral triple defines is a meromorphic function, the so-called zeta function of the triple. The pole structure of the zeta-function defines various dimensions (dimension spectrum). In the commutative case these dimensions can be compared with Hausdorff or Minkowski dimensions which arise if the space is equipped with a metric. A third object is a particular state on the algebra, the spectral state, which in the commutative case corresponds to a measure, the spectral measure. Hence one does not only have the elements of a differential calculus at hand – via the commutator with D – but also an integral. This allows one to define the fourth object, a quadratic form which may be interpreted as a Laplace operator. The fifth and final object we consider here is a K-homology class, an element of noncommutative topology which ought to be of fundamental importance for the underlying geometry. Many more structures can be obtained from the spectral triple of an algebra, but we will not consider them here. Instead we refer to two rather comprehensive books on the subject [12, 22].

How does one construct spectral triples for the commutative algebra of a tiling? The basic idea, which goes back to Christiansen & Ivan, is based on the pair spectral triple. A space[1] consisting of just two points allows for an essentially unique non-trivial spectral triple depending on one parameter only which can be interpreted as the distance between the two points. This is the so-called pair spectral triple. Now one approximates the space by a countable dense subset and declares in a hierarchical way which points are paired up to make a pair spectral triple. The actual spectral triple is then the direct sum of all pair spectral triples. There are of course a lot of choices to make along the way like, for instance, the parameter for the pair spectral triple. One of the surprising aspects of the construction is that it pays off to bundle up certain choices in a so-called choice function and to consider a whole family of spectral triples parametrised by the choice functions. Objects as the ones mentioned above and certain quantities

[1]Here and in the following we mean by a spectral triple for a space a spectral triple for the (commutative) algebra of continuous functions over the space.

derived from them can then be obtained either by taking extremal values or by averaging over the choice functions.

In different disguises this approach has been considered for general compact metric spaces [13, 48], ultrametric Cantor sets [50], fractals [23, 24] and the spaces of tilings and subshifts. With the exception of the results on fractals we review these works in the unifying framework of approximating graphs with an emphasis on tiling and subshift spaces. More precisely, we consider in Section 3 the *spectral triple of a subshift* (Section 3.1), the *ordinary transverse spectral triple of a tiling* (Section 3.2), the *transverse substitution spectral triple* of a substitution tiling (Section 3.3.1), the *longitudinal substitution spectral triple* of a prototile of a substitution tiling (Section 3.3.2), and, by combining the last two, the *full substitution spectral triple* of a substitution tiling (Section 3.3.3). The following sections are then devoted to the study of the above-mentioned objects which can be defined from the data of a spectral triple.

Section 4 is devoted to the zeta-function and the relation between its poles, various dimensions defined for metric spaces, and, most importantly, complexity exponents of subshifts and tilings.

In Section 5 we review the results on Dirichlet forms and Laplacians for ultrametric Cantor sets, discrete hulls of tilings and the continuous hull of the substitution tiling. The first proposal of a Dirichlet form for the spectral triple of an ultrametric Cantor set has been made by Pearson and Bellissard [50]. We discuss it in Section 5.1. It is in some sense not the canonical Dirichlet form, as it involves an averaging over the choice functions, and integration is not defined w.r.t. the spectral state but rather with the ordinary operator trace. It can be computed quite explicitly in the case of transversals of substitution tilings. The significance of this so-called Pearson–Bellissard Laplacian has still to be understood. Given that the abstract transversal of a tiling is a Cantor set which can be provided with an ultrametric the Pearson–Bellissard Laplacian is expected to be related to the motion of the transverse degrees of freedom of the aperiodic medium described by the tiling. But a detailed understanding of this expectation has still to be found.

We then review the results of a Dirichlet form defined for the full substitution spectral triple of a Pisot substitution tiling (Section 5.2). This time integration is defined using the spectral state. The main result is quite different from the Pearson–Bellissard Laplacian. In fact, the Laplacian for the full substitution spectral triple can be interpreted as an elliptic differential operator with constant coefficients defined on the dual group of the group of topological eigenvalues of the dynamical system of the tiling.

Section 6 is devoted to a somewhat unexpected development which has to do with the fact that the spectral triple of a subshift is parametrised by a choice function. As a consequence we have a whole family of distance functions depending on that parameter. One may ask whether these are all equivalent in the sense of Lipschitz continuity. We consider hence the extremal values of the distance function, $\underline{d}$ and $\overline{d}$ and ask when there exists $c > 0$ such that $c^{-1}\underline{d} \leq \overline{d} \leq c\underline{d}$. This turns out to be equivalent to a combinatorial property of the subshift. Such

a constant $c > 0$ exists if and only if the subshift has bounded powers, i.e., there exists $N > 0$ such that no sequence ever contains an N-fold repetition of a word. For a Sturmian subshift of slope θ this combinatorial property is equivalent to the number theoretic property that θ has a bounded continued fraction expansion. This is an example in which noncommutative geometry can say something about a certain combinatorial property of the subshift. This combinatorial property in turn can be understood as a criterion for high aperiodic order of the subshift.

The final section is about the K-homology of compact ultrametric spaces. Any spectral triple defines a K-homology class and therefore via Connes pairing a group homomorphisms on K-theory with values in $\mathbb{Z}$. We consider these group homomorphisms in the context of the spectral triples we construct for compact ultrametric spaces. The flexibility of our construction allows us to design for every such homomorphism a spectral triple which defines it.

2. Spectral triples

2.1. General definition

Definition 2.1. A spectral triple $(A, D, \mathfrak{H})$ is given by a complex involutive associative unital algebra A which is faithfully represented by bounded operators on some Hilbert space $\mathfrak{H}$ together with a self-adjoint operator D on $\mathfrak{H}$ of compact resolvent such that all commutators $[D, a]$, $a \in A$, extend to bounded operators. The spectral triple is called *even* if there exists a grading operator on the Hilbert space such that A is represented by even operators and D is an odd operator.

The basic idea is that $[D, a]$ is the derivative of a thus furnishing a differential calculus on A, although it is a priori not required that $[D, a]$ lies in the image of the representation (and hence defines an element of A).

Given a C^*-algebra A, a spectral triple for A is a spectral triple $(A_0, D, \mathfrak{H})$ in the above sense where A_0 is a dense subalgebra of A (we suppose that the representation of A_0 on $\mathfrak{H}$ is continuous and hence also A is represented on $\mathfrak{H}$). If it is useful, we emphasize the representation π.

There are a couple of additional requirements made usually to ensure good properties and rigidify the theory. These are most often motivated by Riemannian geometry. We will mention in our discussion some of those, keeping an open mind not to constrain to much as our application in mind is to tilings.

We will now quickly browse through the main objects which can be defined by means of a spectral triple and which are or relevance for us.

Connes distance. The formula

$$d_C(\sigma, \sigma') = \sup\{|\sigma(a) - \sigma'(a)| : a \in A, \|[D, a]\| \leq 1\} \tag{1}$$

defines a pseudo-metric on the state space $S(A)$ of A. This pseudo-metric d_C is a metric whenever the representation π is non-degenerate and $A'_D := \{a \in A : [D, a] = 0\}$ is one-dimensional, that is, contains only multiples of the identity.

The state space comes with a natural topology, namely the weak-$*$ topology. One may ask, when is the metric compatible with the weak-* topology? Rieffel has provided a complete answer to this question [54]: assuming that d_C is a metric, it generates the weak-* topology on $S(A)$ if and only if the image of $B_1 := \{a \in A : \|[D,a]\| \leq 1\}$ in the quotient space A/A'_D is pre-compact. While complete, this characterisation is not always easy to verify and we will indeed use more direct methods to verify whether d_C generates the topology or not.

This is already very interesting if A is commutative. In this case A is isomorphic to $C(X)$, for some compact topological space X which is homeomorphic to the closed subset of pure states on A. Eq. (1) restricted to the pure states then becomes

$$d_C(x,x') = \sup\{|f(x) - f(x')| : f \in C(X), \|[D,f]\| \leq 1\}$$

and under Rieffel's conditions d_C generates the topology of X. It is therefore quite natural to require that Rieffel's conditions are satisfied. On the other hand, however, if the construction of the spectral triple follows a natural principle one can use the criterion of whether or not the Connes distance generates the topology as a characterisation of the algebra A, or the space X; this is the basis of the characterisation of order we exhibit in Section 6.

Zeta function. Since the resolvent of D is supposed compact $\mathrm{Tr}(|D|^{-s})$ can be expressed as a Dirichlet series in terms of the eigenvalues of $|D|$.[2] The spectral triple is called *finitely summable* if the Dirichlet series is summable for some $s \in \mathbb{R}$ and hence defines a function

$$\zeta(z) = \mathrm{Tr}(|D|^{-z}),$$

on some half-plane $\{z \in \mathbb{C} : \Re(z) > s_0\}$ which is called the *zeta-function* of the spectral triple. Under the right circumstances, ζ admits a meromorphic extension to the whole complex plane and then its pole structure yields interesting information. We will see that in particular the smallest possible value for s_0 in the above (the *abscissa of convergence* of the Dirichlet series) that is, the largest pole on the real axis, is related to the complexity of the tiling. This number s_0 is also called the *metric dimension* of the spectral triple.

Spectral state and measure. Under the right circumstances, the limit

$$\mathcal{T}(T) = \lim_{s \to s_0^+} \frac{1}{\zeta(s)} \mathrm{Tr}(|D|^{-s}T).$$

exists for a suitable class of operators on $\mathfrak{H}$. It then defines a positive linear functional on this class of operators which we call the *spectral state* defined by $(A, D, \mathfrak{H})$. It is of particular interest already if $A = C(X)$ is commutative. Then the restriction of $\mathcal{T}$ to A defines by the Riesz representation theorem a measure on X. We call that measure the *spectral measure* defined by $(A, D, \mathfrak{H})$.

[2] For simplicity we suppose (as will be the case in our applications) that $\ker(D)$ is trivial, otherwise we would have to work with $\mathrm{Tr}_{\ker(D)^\perp}(|D|^{-s})$ or remove the kernel of D by adding a finite rank perturbation.

Dirichlet forms. Expressions of the type $(a,b) \mapsto \mathcal{T}([D,\pi(a)]^*[D,\pi(b)]$, where $\mathcal{T}$ is a state, define quadratic forms on suitable subspaces of A which, under the right circumstances, can be extended to Dirichlet forms on the L^2-space over A w.r.t. to the spectral state. In this definition $\mathcal{T}$ is often, but not always also the spectral state. Indeed, we will see that the choice $\mathcal{T} = \mathrm{Tr}$ yields interesting Dirichlet forms.

The interest in Dirichlet forms comes from the fact that they define Markov processes with generators which in the classical cases correspond to Laplace–Beltrami operators. Under the right circumstances, a spectral triple provides us therefore with a Laplacian.

***K*-homology.** The spectral triple yields directly an unbounded Fredholm module and therefore the representative of a K-homology class of A. By the Connes pairing of K-homology with K-theory the spectral triple defines thus a functional on $K_*(A)$.

The points mentioned above are also interesting for commutative algebras. We will in particular consider here the cases $A = C(\Xi)$ or $A = C(\Omega)$ where Ξ is the discrete and Ω the continuous hull of a tiling.

2.2. Spectral triples for metric spaces

Let X be a compact topological Hausdorff space. There exist various constructions for spectral triples for X, *i.e.*, for the algebra $C(X)$. These constructions are designed to fulfill additional properties. For instance, if X is a Riemannian spin manifold with metric g then one can use the Hilbert space of L^2-spinors with its standard Dirac operator defined with the help of the spin structure to obtain a spectral triple which has the property that the Connes' metric is equal to the one determined by g. We won't describe this construction here, as it draws heavily on differential geometry and we are interested in spaces which are far from being a manifold. We would like to discuss two properties that spectral triples can satisfy – though not necessarily at the same time. One is that the associated Connes' metric is equivalent to the original metric with a Lipschitz constant which is arbitrarily close to 1, and the other is that the metric dimension coincides with the Hausdorff dimension or with the lower box counting dimension of X.

The simplest case of an even spectral triple seems to us the spectral triple of a pair of points. X consists of two points so the algebra is $C(X) = \mathbb{C}^2$ acting diagonally on $\mathfrak{H} = \mathbb{C}^2$ and the Dirac operator is $D = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Its associated Connes distance gives the two points distance 1. Multiplying D by $\frac{1}{d}$ the two points get distance d. This spectral triple is the base for all what follows. The idea is to approximate the metric space by finite point sets and to encode that two points are considered to be neighbours with the help of (horizontal) edges linking the two points. The spectral triple of the approximation will then be a (finite) direct sum of pair spectral triples. Next, the approximation is refined, *i.e.*, X is approximated by more points which eventually become dense. Taking a sequence of finer and finer approximations we will end up with a spectral triple for the space which will

be a countable direct sum of pair spectral triples. This idea occurs in the work of Christensen and Ivan [13]. In the next section we will make this idea precise.

2.2.1. Rooted trees and approximating graphs. Rooted trees are the first (and most fundamental) ingredient of our construction of an approximating graph and its associated spectral triple. A rooted tree is a connected graph $\mathcal{T}$ without cycles (as an un-oriented graph) which has a distinguished vertex called the root. We denote by $\mathcal{T}^{(0)}$ the vertices and by $\mathcal{T}^{(1)}$ the edges of the tree. The edges of the graph may be oriented in such a way that they point away from the root. By a path on the tree we then mean a sequence of edges such that the endpoint (or range) of the nth edge corresponds to the startpoint (or source) of the $n+1$th edge. Here the start and endpoint are defined w.r.t. the orientation. We write $v \preceq v'$ if there is a path (possibly of zero length) from v to v'. Any vertex is the endpoint of a unique path which starts at the root vertex. We denote by $\mathcal{T}_n^{(0)}$ the vertices whose corresponding unique path has length n and call n also the level of the vertex. In particular $\mathcal{T}_0^{(0)}$ contains only the root vertex. Any vertex besides the root has one incoming edge. A branching vertex is a vertex which has at least two outgoing edges. We assume that it has finitely many outgoing edges and hence that each $\mathcal{T}_n^{(0)}$ is finite. For a vertex $v \in \mathcal{T}_n^{(0)}$ we let $\mathcal{T}^{(0)}(v) := \{v' \in \mathcal{T}_{n+1}^{(0)} : v \preceq v'\}$ be the set of its successors.

The boundary of the tree, denoted $\partial\mathcal{T}$ is defined as the set of infinite paths on $\mathcal{T}$ which start at the root vertex. We equip it with the topology whose basis is in one-to-one correspondence with the vertices of $\mathcal{T}$: each vertex v defines a set $[v]$ of that basis, namely $[v]$ is the set of all infinite paths on $\mathcal{T}$ which pass through v. If $v \in \mathcal{T}_n^{(0)}$ then the complement of $[v]$ is the union of all $[w]$ where $v \neq w \in \mathcal{T}_n^{(0)}$. Hence $\partial\mathcal{T}$ is totally disconnected space (it has a basis of clopen sets). Moreover it is metrizable. Let ξ, η be two infinite paths on $\mathcal{T}$ which start at root. If they are distinct then they will branch at a certain vertex, or stated differently, they agree on a finite path from root on. Let $\xi \wedge \eta$ be that finite path and $|\xi \wedge \eta|$ its (combinatorial) length. Then, given any strictly decreasing function $\delta : \mathbb{R}^+ \to \mathbb{R}^+$ which tends to 0 at $+\infty$,

$$d(\xi, \eta) = \delta(|\xi \wedge \eta|)$$

defines a metric which induces the topology above. d is in fact an ultrametric, that is, it satisfies $d(x, y) \leq \max\{d(x, z), d(y, z)\}$ for all $x, y, z \in X$. Our assumption that $\mathcal{T}_n^{(0)}$ is finite implies that $\partial\mathcal{T}$ is compact.

The data of an approximating graph for a compact space X are the following.

1. A rooted tree $\mathcal{T} = (\mathcal{T}^{(0)}, \mathcal{T}^{(1)})$. We assume that $\partial\mathcal{T}$ is a dense $1:1$ extension of X, that is, there is a continuous surjection $q : \partial\mathcal{T} \to X$ such that the points which have a unique preimage form a dense set.
2. A non-empty symmetric subset

$$\mathcal{H} \subset \{(v, v') \in \mathcal{T}^{(0)} \times \mathcal{T}^{(0)} : v \not\preceq v', v' \not\preceq v\}$$

whose elements we call horizontal edges interpreting them as the edges of a graph $(\mathcal{V}, \mathcal{H})$. Here $\mathcal{V} \subset \mathcal{T}^{(0)}$ is the subset of vertices which are the source (and hence also the range) of at least one element in $\mathcal{H}$. This graph, which we call the horizontal graph, has no loop edges and no multiple edges and for each edge $(v, v') \in \mathcal{H}$ contains also the edge with opposite orientation (v', v). We suppose that $(\mathcal{V}, \mathcal{H})$ is locally finite in the sense that each vertex $v \in \mathcal{V}$ is connected only to finitely many incoming or outgoing edges.

3. There is a so-called length function $\delta : \mathcal{H} \to \mathbb{R}^{>0}$ which satisfies $\delta(v', v) = \delta(v, v')$ and two further conditions: (i) for all $\epsilon > 0$ the set $\{h \in \mathcal{H} : \delta(h) > \epsilon\}$ is finite (ii) there is a strictly decreasing sequence $(\delta_n)_n$ tending to 0 at ∞ such that $\delta(v, v') \geq \delta_{|v \wedge v'|}$.
4. There is a so-called choice function $\tau : \mathcal{T}^{(0)} \to \partial\mathcal{T}$ which satisfies
   (a) $\tau(v)$ goes through v.
   (b) if $w \prec v$ then $\tau(w) = \tau(v)$ if and only if $\tau(w)$ passes through the vertex v.
   (c) q is injective on the image of τ.

   Note that we have the following sequences of maps

$$\mathcal{H} \underset{s}{\overset{r}{\rightrightarrows}} \mathcal{V} \overset{\tau}{\to} \partial\mathcal{T} \overset{q}{\to} X \tag{2}$$

where $r(v, v') = v'$ and $s(v, v') = v$. Since $q \circ \tau \times q \circ \tau : \mathcal{H} \to X \times X$ is injective the length function defines a distance function on $q(\tau(\mathcal{V}))$.

Definition 2.2. The approximating graph of the above data (1)–(4) is the (metric) graph G_τ with vertices $V = q(\tau(\mathcal{V}))$ and edges $E = \{\big(q \circ \tau(v), q \circ \tau(v')\big) : (v, v') \in \mathcal{H}\}$ equipped with the length function $\big(q \circ \tau(v), q \circ \tau(v')\big) \mapsto \delta(v, v')$.

For each pair of vertices v, v' which are linked by an edge in $\mathcal{H}$ we choose (arbitrarily) an order calling the edge defined by the vertices in that order positively oriented and the one with the reversed order negatively oriented. This splits $\mathcal{H}$ into two parts $\mathcal{H}^{\pm}$ according to the chosen orientation of the edges. The splitting is, of course, non-canonical but the choice of orientation will not affect the final results. Since E is in bijective correspondence with $\mathcal{H}$ this orientation carries over to E. We denote by $\cdot^{\mathrm{op}}$ the operation of changing the orientation of an edge. Conditions (b) and (c) above imply that $q \circ \tau(v) \neq q \circ \tau(v')$ if $(v, v') \in \mathcal{H}$. Let

$$C(\partial\mathcal{T})_0 = \left\{ f \in C(\partial\mathcal{T}) : \sup_{(v,v') \in \mathcal{H}} \frac{|f(\tau(v)) - f(\tau(v'))|}{\delta(v, v')} < \infty \right\}. \tag{3}$$

Lemma 2.3. *$C(\partial\mathcal{T})_0$ is a dense subalgebra of $C(\partial\mathcal{T})$.*

Proof. Let χ_v denote the characteristic function on $[v]$. Since $[v]$ is clopen $\chi_v \in C(\partial\mathcal{T})$. We claim that the expression $\chi_v(\tau(v_2)) - \chi_v(\tau(v_1))$ can be non-zero only if $v_1 \wedge v_2 \prec v$. Indeed, if this is not the case then either $v_1 \wedge v_2 \succeq v$ or $v_1 \wedge v_2 \wedge v \prec v, v_1 \wedge v_2$. In the first case $\tau(v_1)$ and $\tau(v_2)$ both contain v and in the second both

do not. Since $v_1 \wedge v_2 \prec v$ implies $\delta(v_1, v_2) \geq \delta_{|v|}$ we have

$$\frac{|\chi_v(\tau(v_2)) - \chi_v(\tau(v_1))|}{\delta(v_1, v_2)} \leq \frac{2}{\delta_{|v|}}$$

and so we see that $\chi_v \in C(\partial\mathcal{T})_0$. Moreover, since $\partial\mathcal{T}$ is totally disconnected the algebra generated by characteristic functions on $[v]$ is dense in $C(\partial\mathcal{T})$. □

2.2.2. The spectral triple of an approximating graph. We first consider a spectral triple over $\partial\mathcal{T}$. It depends on the above data except the surjection q.

Theorem 2.4. *Consider a rooted tree $\mathcal{T}$ with a set of horizontal edges $\mathcal{H}$, a length function δ, and a choice function τ as above. Let $\mathfrak{H} = \ell^2(\mathcal{H})$ and represent $C(\partial\mathcal{T})$ on $\mathfrak{H}$ by π_τ,*

$$\pi_\tau(f)\psi(h) = f(\tau(s(h)))\psi(h).$$

Let the Dirac operator be given by

$$D\psi(h) = \frac{1}{\delta(h)}\psi(h^{op}).$$

Then $(C(\partial\mathcal{T}), D, (\mathfrak{H}, \pi_\tau))$ defines an even spectral triple w.r.t. the decomposition $\mathfrak{H}^\pm = \ell^2(\mathcal{H}^\pm)$ defined by the orientation of the edges.

Proof. Since $|D|^{-1}\psi(h) = \delta(h)\psi(h)$ we see that D has compact resolvent if and only if for all $\epsilon > 0$ the set $\{h \in \mathcal{H} : \delta(h) > \epsilon\}$ is finite.

Furthermore

$$[D, \pi_\tau(f)]\psi(v_1, v_2) = \frac{1}{\delta(v_1, v_2)}\big(f(\tau(v_2)) - f(\tau(v_1))\big)\psi(v_2, v_1).$$

Hence $C(\partial\mathcal{T})_0$ is precisely the subalgebra of functions f for which the commutator $[D, \pi_\tau(f)]$ is bounded. □

Using the surjection q we can define a spectral triple on X.

Theorem 2.5. *Consider an approximating graph G_τ as above with length function δ. Let $\mathfrak{H} = \ell^2(E)$ and represent $C(X)$ on $\mathfrak{H}$ by π,*

$$\pi(f)\psi(e) = f(s(e))\psi(e).$$

Let the Dirac operator be given by

$$D\psi(e) = \frac{1}{\delta(e)}\psi(e^{op}).$$

If $q^{*-1}(C(\partial\mathcal{T})_0) \subset C(X)$ is dense in $C(X)$ then $(C(X), D, \mathfrak{H})$ defines an even spectral triple w.r.t. the decomposition $\mathfrak{H}^\pm = \ell^2(E^\pm)$ defined by the orientation of the edges.

Proof. The only issue is the question whether $C(X)_0 = \{f \in C(X) : \|[D, \pi(f)]\| < \infty\}$ is dense in $C(X)$. By construction $q^{*-1}(C(\partial\mathcal{T})_0) \subset C(X)_0$ and so our hypothesis guarantees that property. □

We call this spectral triple the *spectral triple of the approximation graph* (G_τ, δ). Note that if q is a bijection and hence X homeomorphic to $\partial\mathcal{T}$ then the two spectral triples are the same. In the first formulation the dependence on the choice function shows up in the definition of the representation whereas in the second formulation it is the embedding of the graph in X which depends on it.

Without any further conditions on δ nothing can be said about whether the Connes distance induces the topology of X.

How do such triples arise? We will discuss in the next sections examples with canonical tree structure, one being the case of a compact ultrametric space and the other being the case of a substitution tiling space. In these cases, more or less natural choices for the horizontal edges $\mathcal{H}$ and the length function δ can be argued for. This is not the case for the choice function τ which therefore has to be regarded as a parameter. In [50] it is interpreted as the analogue of a tangent vector of a manifold.

An arbitrary compact metric space (X, d) can be described by an approximation graph as above, although the involved graph is not canonical. Palmer [48] starts with a sequence $(\mathcal{U}_n)_n$ of open covers of X. We call the sequence *refined* if $\mathcal{U}_{n+1}$ is finer than $\mathcal{U}_n$ for all n, and *resolving* if it is refined and $\operatorname{diam}\mathcal{U}_n \overset{n\to\infty}{\longrightarrow} 0$.[3] A resolving sequence *separates points*: for any $x, y \in X$, $x \neq y$, there exists n and $U, U' \in \mathcal{U}_n$ such that $x \in U$, $y \in U'$, and $U \cap U' = \emptyset$. In general such a sequence $(\mathcal{U}_n)_n$ is by no means unique. There is a graph $\mathcal{T} = (\mathcal{T}^{(0)}, \mathcal{T}^{(1)})$ associated with such a sequence: vertices in $\mathcal{T}_n^{(0)}$ are in one-to-one correspondance with open sets in $\mathcal{U}_n$; this correspondance is written $v \leftrightarrow U_v$. The inclusion $U_v \subset U_w$ for $U_v \in \mathcal{U}_n$, $U_w \in \mathcal{U}_{n-1}$ defines an edge of $\mathcal{T}_n^{(1)}$ with source $w \in \mathcal{T}_{n-1}^{(0)}$ and range $v \in \mathcal{T}_n^{(0)}$. An infinite path $\xi \in \partial\mathcal{T}$ stands for a sequence of open sets $(U_{\xi_n})_n$ such that $U_{\xi_n} \subset U_{\xi_{n+1}}$. As the sequence $(\mathcal{U}_n)_n$ is resolving, the intersection $\bigcap_n U_{\xi_n}$ contains a single point $q(\xi)$. This defines a map $q : \partial\mathcal{T} \to X$, which turns out to be a continuous surjection. One has now to introduce choice and length functions to build an approximation graph as above and obtain the spectral triple of Theorem 2.5.

Palmer considers choice functions which are slightly more general than here (sets of horizontal edges $\mathcal{H}$ can be deduced from his choice functions). The length function δ is simply the actual distance of the points in X. These choices fix an approximation graph (G_τ, δ). Palmer's main concern is to show the existence of a resolving sequence $(\mathcal{U}_n)_n$ such that, for any choice function, he gets a spectral triple recovering the Hausdorff dimension of X and the Hausdorff measure on Borel sets of X.

Theorem 2.6 ([48]). *There exist a resolving sequence $(\mathcal{U}_n)_n$ such that, for any choice function, the spectral triple for X defined by the approximation graph constructed from the above data has*

1. *metric dimension s_0 equal to the Hausdorff dimension of X, and*
2. *spectral measure equal to the (normalized) Hausdorff measure on X.*

[3] The diameter of a covering $\mathcal{U}$, written $\operatorname{diam}\mathcal{U}$, is the supremum of the diameters of the sets of $\mathcal{U}$.

As a corollary, the Hausdorff measure of X can be shown to be

$$\lim_{n\to\infty} \sum\nolimits_{U\in\mathcal{U}_n} (\operatorname{diam} U)^{s_0}.$$

Finally we mention that the approach of Christensen & Ivan [13] (which predates [48]) can be recast as well in the framework of approximating graphs. One does not have to work with an abstract tree and choice functions to construct such a graph but may simply proceed as follows: Start with a sequence $(V_n)_n$ of finite subsets V_n of X such that $V_n \subset V_{n+1}$ and their union $\bigcup_n V_n$ is dense in X. For each n choose a non-empty symmetric subset $E_n \subset V_n \times V_n$ and define $\delta(x,y) = d(x,y)$, *i.e.*, the length of edge (x,y) corresponds to their distance in X. Now one has directly the approximating graph (V,E) with metric. But there is just enough structure to recover the metric aspects, namely one gets:

Theorem 2.7 ([13]). *Let (X,d) be a compact metric space. For any constant $\alpha > 1$ there exists a sequence $(V_n)_n$ of finite point sets $V_n \subset X$ together with a choice of horizontal pairs E_n as above, such that the spectral triple of the approximating graph (V,E) with length function δ yields a spectral metric d_C which satisfies*

$$d(x,y) \le d_C(x,y) \le \alpha d(x,y).$$

2.3. Spectral triples for compact ultrametric spaces

In this section we consider spectral triples for compact ultrametric spaces. With a particular choice for the horizontal egdes, we obtain the spectral triples of Pearson–Bellissard [50]. The general case has been discussed in [39].

A compact ultrametric space (X,d) is a metric space for which the metric satisfies the strong triangle inequality

$$d(x,y) \le \max\{d(x,z), d(y,z)\} \tag{4}$$

for all $x,y,z \in X$. We suppose that X has infinitely many points (the finite case being simpler). Such spaces arise as subshifts or as transversals of spaces ("discrete tiling spaces") of non-periodic tilings with finite local complexity. The property (4) implies that the open δ-balls $B_\delta(x) = \{y \in X : d(x,y) < \delta\}$ satisfy either $B_\delta(x) = B_\delta(y)$ or $B_\delta(x) \cap B_\delta(y) = \emptyset$. In particular there is a unique cover by δ-balls (δ-cover). Moreover this cover is a partition and hence X is totally disconnected. Furthermore, the image of d contains exactly one accumulation point, namely 0. In other words, there exists a strictly decreasing sequence $(\delta_n)_n$ converging to zero such that $\operatorname{im} d = \{\delta_n : n \in \mathbb{N}\}$. If we take $\mathcal{U}_n$ to be the δ_n-cover of X we obtain a canonical refined resolving sequence of coverings. Its associated tree $\mathcal{T}$ is the so-called Michon tree [44, 50, 48]. In particular, the vertices of level n of $\mathcal{T}$ correspond to the clopen δ_n-balls of the δ_n-cover, and the root corresponds to all of X. The sets of the covering are also closed, so given an infinite path ξ, the sequence of vertices through which it passes defines a nested sequence of compact sets whose radius tends to 0. It defines a point $x(\xi) \in X$. The map $\xi \mapsto x(\xi)$ is even injective in the case of ultrametric spaces and thus furnishes a homeomorphism between $\partial\mathcal{T}$ and X.

Now that we have the canonical tree associated with the compact ultrametric space we need to choose horizontal edges. Any vertex has one incoming vertical

edge. A branching vertex is a vertex which has at least two outgoing (vertical) edges. For a vertex $v \in \mathcal{T}_n^{(0)}$ let $\mathcal{T}^{(0)}(v) := \{v' \in \mathcal{T}_{n+1}^{(0)} : v \preceq v'\}$. A canonical choice for the horizontal edges $\mathcal{H}$ is to introduce an edge between any pair of distinct vertices of $\mathcal{T}^{(0)}(v)$, and this for any branching vertex. This is the *maximal* choice. A *minimal* choice would be to choose, for each branching vertex v, two distinct vertices of $\mathcal{T}^{(0)}(v)$ and to introduce two horizontal edges, one for each direction, only for these two. This case has been considered in [50]. Note that in both choices we obtain a grading for the horizontal edges: In the maximal choice $\mathcal{H}^{\max} = \bigcup_n \mathcal{H}_n^{\max}$ with

$$\mathcal{H}_n^{\max} = \{(v', v'') \in \mathcal{T}^{(0)}(v), v \in \mathcal{T}_{n-1}^{(0)}, v' \neq v''\}$$

and a minimal choice is a subset $\mathcal{H}^{\min} = \bigcup_n \mathcal{H}_n^{\min}$, with $\mathcal{H}_n^{\min} \subset \mathcal{H}_n^{\max}$. More generally we may consider horizontal edges of the form

$$\mathcal{H} = \bigcup_n \mathcal{H}_n, \qquad \mathcal{H}_n \subset \mathcal{H}_n^{\max}. \tag{5}$$

The natural length function for all these cases is the function determined by the radii of the balls, *i.e.*, $\delta(v, v') = \delta_n$, where $(v, v') \in \mathcal{H}_n$.

Any choice function τ provides us with a spectral triple for the compact ultrametric space X according to Theorems 2.4 or 2.5 (since q is a homeomorphism the two theorems are equivalent in the present case).

Definition 2.8. By a spectral triple for a compact ultrametric space we mean the spectral triple given by the above data: its Michon tree, a choice of horizontal edges $\mathcal{H}$ as in (5), length function determined by the radii of the balls, and any choice function. If $\mathcal{H} = \mathcal{H}^{\min}$ we refer to the spectral triple also as Pearson–Bellissard spectral triple.

Lemma 2.9. *Consider a spectral triple for a compact ultrametric space. Its spectral distance d_C bounds the original metric $d(x, y) \leq d_C(x, y)$. If, for any $v', v'' \in \mathcal{T}^{(0)}(v)$ (with $v \in \mathcal{T}^{(0)}$), there is a path of edges in $\mathcal{H}$ linking v with v' then G_τ is connected and hence d_C a metric.*

The condition on $\mathcal{H}$ formulated in the lemma is obviously satisfied for the maximal choice $\mathcal{H}^{\max}$.

3. Concrete examples for subshifts and tilings

We discuss applications of the previous constructions to subshifts and to tilings of finite local complexity.

Example 3.1. We will illustrate some of our constructions using the example of the Fibonacci tiling. We will consider three versions of it:

1. The Fibonacci one-sided subshift;
2. The Fibonacci two-sided subshift;
3. The one-dimensional Fibonacci substitution tiling.

The Fibonacci subshift (one-sided or two-sided) is a subshift of the full (one-sided or two-sided) shift over the alphabet $\mathcal{A} = \{a, b\}$ defined usually by the substitution σ, $\sigma(a) = ab$, $\sigma(b) = a$. This means by definition that the elements of the subshift are infinite sequences of letters, that is, functions $\mathbb{N} \to \mathcal{A}$ (one-sided) or functions $\mathbb{Z} \to \mathcal{A}$ (two-sided), whose finite parts (words) are allowed in the sense that they occur as subwords of $\sigma^n(a)$ for some n (which depends on the word). Note that instead of σ we could also take σ^2 or even the substitution $a \mapsto baa$, $b \mapsto ba$ to define the same subshift, because all these substitutions yield the same notion of allowed words. The advantage of the latter substitution is that it manifestly forces its border.

The one-dimensional Fibonacci substitution tiling is the suspension of the two-sided subshift, in which the letters are realised as intervals (one-dimensional tiles), a by an interval of length $(\sqrt{5}+1)/2$ and b by one of length 1. Alternatively this is a canonical cut-and-project tiling of $\mathbb{Z}^2 \subset \mathbb{R}^2$ onto $\mathbb{R}^1$ as the line throughout the origin with irrational slope $(\sqrt{5}+1)/2$, and window the half-open unit cube $[0,1) \times [0,1)$.

3.1. One-sided subshifts and the tree of words

Our first application is to the space of one-sided sequences of a subshift. The construction, which is described in [39, 33], is based on the so-called tree of words of the subshift.

Recall that a one-sided full shift over a finite alphabet $\mathcal{A}$ is the set of sequences $\mathcal{A}^{\mathbb{N}}$ considered as a compact topological space whose topology is that of pointwise convergence. On this space we have an action of $\mathbb{N}$ by left shift, that is, dropping the first letter. A subshift is a closed, shift-invariant subset of the full shift and its language $\mathcal{L}$ is the set of finite words occurring in the sequences of the subshift.

The tree of words for a one-sided subshift with language $\mathcal{L}$ is defined as follows: the vertices of level n, noted $\mathcal{T}_n^{(0)}$, are the (allowed) words of length n and the empty word corresponds to the root. Given a word $w \in \mathcal{L}$ and any of its one-letter extensions $wa \in \mathcal{L}$, $a \in \mathcal{A}$, we draw an edge from w to wa. Hence any word has exactly one incoming edge and at least one outgoing edge. A word is called right-special if it can be extended to the right by one letter in more than one way. A right-special word corresponds thus to a branching vertex in the tree. Note that if the subshift is aperiodic then there is at least one right-special word per length.

We will consider three interesting choices for the horizontal edges.

- *The maximal choice* $\mathcal{H}^{\max}$ as introduced in Section 2.3. With the maximal choice there is a horizontal edge between any two distinct one letter extensions of a right-special word. The level of such an edge is thus equal to the length of the word plus 1.
- *A minimal choice* as introduced in Section 2.3. For each right-special word one chooses a pair of distinct one letter extensions which then are linked by two edges (one for each orientation).

- *The privileged choice* $\mathcal{H}^{\mathrm{pr}}$. This is like the maximal choice but with respect to a certain subset of so-called *privileged words*. Concretely, there is a horizontal edge between any two distinct privileged extensions of a privileged word. The definition of privileged words is based on return words. A word $v \in \mathcal{L}$ is a complete first return to a word $u \in \mathcal{L}$ if u occurs exactly twice in v, namely once as a prefix and once as a suffix. By convention, a letter is a complete first return to the empty word. A privileged word is defined iteratively: the empty word is a privileged word, and a complete first return to a privileged word is a privileged word. Now $(v, v') \in \mathcal{H}^{\mathrm{pr}}$ if and only if v and v' are privileged and there is a privileged word $u \in \mathcal{T}^{(0)}$ such that v and v' are distinct complete first returns to u.

We will specify the length function later according to our needs.

Definition 3.2. By a spectral triple for a one-sided subshift we mean the spectral triple as defined in Theorems 2.4 or 2.5 given by the above data: the tree of words, a choice of horizontal edges $\mathcal{H}$ as above, a choice of length function δ, and any choice function.

We are interested in two different types of subshifts: minimal aperiodic subshifts in which case the subshift (or its two-sided version) stands for the symbolic version of a one-dimensional tiling, or subshifts of finite type, which arise in the context of substitution tilings.

Example. The following figure shows four levels of the tree of words for the one-sided Fibonacci subshift of Example 3.1 together with horizontal edges of two different types. As it has a unique right-special factor per length, each vertex of the tree has at most two successors and therefore the minimal and the maximal choice for the horizontal edges coincide.

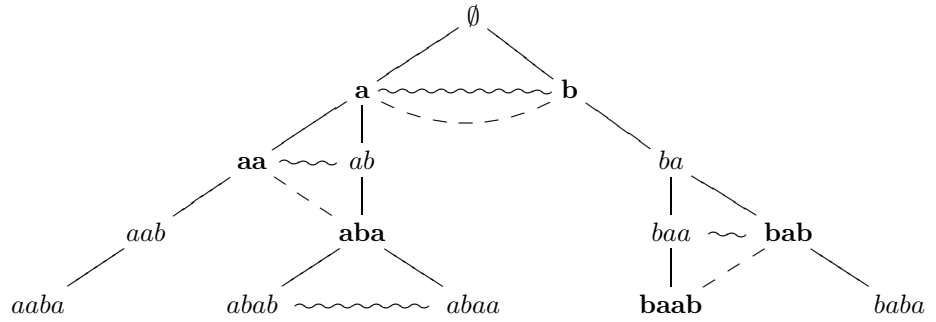

FIGURE 1. Four levels of the tree of words $\mathcal{T}$ for the one-sided Fibonacci subshift. Letters stand for the vertices $\mathcal{T}^{(0)}$ of the tree, and unbroken lines for its edges $\mathcal{T}^{(1)}$. Bold letters stand for privileged words. Drawn are also the horizontal edges for the maximal choice $\mathcal{H}^{\max}$ (curvy lines) and the horizontal edges for the privileged words $\mathcal{H}^{\mathrm{pr}}$ (broken lines).

3.1.1. One-sided subshifts of finite type. Consider a finite oriented graph $\mathcal{G} = (\mathcal{G}^{(0)}, \mathcal{G}^{(1)})$ such as the graph associated with a substitution defined on an alphabet. It defines a one-sided subshift of finite type: this is the subshift whose alphabet is $\mathcal{G}^{(0)}$ and whose language is given by the finite paths on $\mathcal{G}$. The tree of words associated with the subshift of finite type looks as follows: $\mathcal{T}_0^{(0)}$ contains the root vertex, $\mathcal{T}_1^{(0)} = \mathcal{G}^{(0)}$, and $\mathcal{T}_{n+1}^{(0)} = \Pi_n(\mathcal{G})$ is the set of paths of length n on $\mathcal{G}$. Furthermore, $\mathcal{T}^{(1)}$ contains one edge joining the root vertex to each $v \in \mathcal{T}_1^{(0)}$. It contains also, for each path γ over $\mathcal{G}$ and each edge $\epsilon \in \mathcal{G}^{(1)}$ with $s(\epsilon) = r(\gamma)$, one edge joining γ to $\gamma\epsilon$ (we denote the latter edge by $(\gamma, \gamma\epsilon)$). It is clear that $\partial\mathcal{T} = \Pi_\infty(\mathcal{G})$, the set of infinite paths over $\mathcal{G}$.

When the tree of words is built from such a graph, it is possible to define a stationary Bratteli diagram, with a somewhat simpler description than the tree, such that the set of paths on the diagram corresponds canonically to the set of paths on the tree. Such a stationary Bratteli diagram exhibits more clearly the underlying self-similarity than the tree. We will not give details for the construction based on Bratteli diagrams, but the construction suggests that it is natural to have a *self-similar* choice of horizontal edges, in the sense that it should only depend on $\mathcal{G}$. Choose a symmetric subset

$$\hat{\mathcal{H}} \subseteq \left\{(\varepsilon, \varepsilon') \in \mathcal{G}^{(1)} \times \mathcal{G}^{(1)} \ : \ \varepsilon \neq \varepsilon',\ s(\varepsilon) = s(\varepsilon')\right\}$$

which we call *fundamental horizontal edges* and then "lift" these to horizontal edges as follows:

$$\mathcal{H}_n = \{(\gamma\epsilon, \gamma\epsilon') : \gamma \in \Pi_n(\mathcal{G}), (\varepsilon, \varepsilon') \in \hat{\mathcal{H}}, r(\gamma) = s(\epsilon)\} \tag{6}$$

Note that vertices of $\mathcal{T}^{(0)}$ are by definitions paths on $\mathcal{G}$, so the equation above defines indeed an horizontal edge as a pair of vertices on the tree. We fix an orientation on $\hat{\mathcal{H}}$ and carry this orientation over to $\mathcal{H}_n$.

We call the length function *self-similar* if there exists a $0 < \rho < 1$ such that for all $h \in \mathcal{H}_n$

$$\delta(h) = \rho^n.$$

The role of the choice function is to associate an infinite extension to each word. This can also be understood in a way that for each word of length n we make a choice of one-letter extension. In the context of subshifts of finite type it is natural to restrict the choice function in the following way. Let us suppose that $\mathcal{G}$ is connected in the stronger sense that for any two vertices v_1, v_2 there exists a path from to v_1 to v_2 and a path from v_2 to v_1, and that it contains a one-edge loop.[4] We fix such a one-edge loop ϵ^*. Consider a function $\hat{\tau} : \mathcal{G}^{(1)} \to \mathcal{G}^{(1)}$ satisfying that for all $\varepsilon \in \mathcal{G}^{(1)}$:

1. if $r(\varepsilon)$ is the vertex of ϵ^* then $\hat{\tau}(\varepsilon) = \epsilon^*$,
2. if $r(\varepsilon)$ is not the vertex of ϵ^* then $\hat{\tau}(\varepsilon)$ is an edge starting at $r(\varepsilon)$ and such that $r(\hat{\tau}(\varepsilon))$ is closer to the vertex of ϵ^* in $\mathcal{G}$.

[4] By going over to a power of the substitution matrix we can always arrange that the substitution graph has these properties if the substitution is primitive.

Then $\hat{\tau}$ defines an embedding of $\Pi_n(\mathcal{G})$ into $\Pi_{n+1}(\mathcal{G})$ by $\varepsilon_1 \cdots \varepsilon_n \mapsto \varepsilon_1 \cdots \varepsilon_n \hat{\tau}(\varepsilon_n)$ and hence, by iteration, into Π_∞. The corresponding inclusion $\Pi_n(\mathcal{G}) \hookrightarrow \Pi_\infty(\mathcal{G})$ is our choice function τ.

Definition 3.3. With the above self similar choice of horizontal, length function, and choice function we call the spectral triple of a subshift of finite type *self-similar*.

Example. Consider the following graph $\mathcal{G}$ which, as we will see further down, is the substitution graph of the substitution underlying Example 3.1.

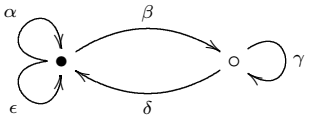

The tree based on this graph (up to level three) is presented in the next figure (due to the lack of space the edges of the last level are partly bundled up). We also draw the (unoriented) horizontal edges arising for the (deliberate) choice $\hat{\mathcal{H}} = \{(\beta, \epsilon), (\gamma, \delta)\}$. For the edge loop we may choose for instance $\epsilon^* = \epsilon$.

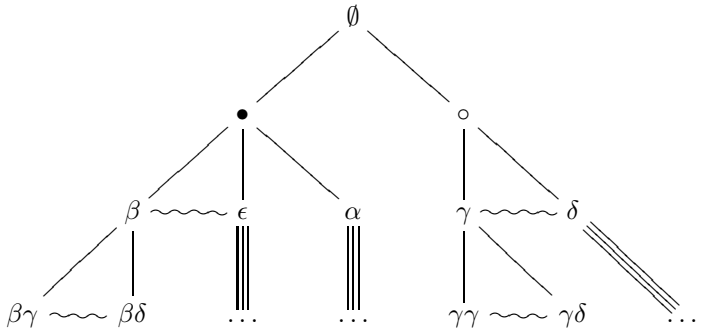

FIGURE 2. Three levels of the tree of words $\mathcal{T}$ for the graph $\mathcal{G}$. Letters stand for the vertices $\mathcal{T}^{(0)}$ of level greater than one, and unbroken lines for the edges $\mathcal{T}^{(1)}$. Horizontal edges are liftings of elements of $\hat{\mathcal{H}}$ and shown by curvy lines.

3.2. The tree of patches and the ordinary transverse spectral triple of a tiling

The construction of the tree of words for one-sided subshift can be generalized to two-sided shifts, $\mathbb{Z}^d$-subshifts, and even tilings of finite local complexity [39]. This generalization is based on the definition of an r-patch. The most common definition in the context of tilings is to pick consistently a privileged point in each tile (puncture), for example their barycenter. The *transversal* of the tiling space is then the set of all tiles which have a puncture at the origin. An r-patch is a patch which has a puncture on 0 and just covers $B_r(0)$, the Euclidean r-ball around the origin.[5] Recall that finite local complexity means that, for any $r > 0$, there

[5]For $\mathbb{Z}^d$ subshifts – which can be viewed as tilings by coloured cubes – it might be more handy to use cubes instead of Euclidean balls.

are only finitely many r-patches. The larger r the more r-patches there are. The number of r-patches is a semi-continuous function of r (the so-called complexity function) and the points where this function jumps form an increasing sequence $(r_n)_n$ of $\mathbb{R}^+$. Given an r_n-patch v, its diameter is noted $|v| := r_n$.

The tree of patches $(\mathcal{T}^{(0)}, \mathcal{T}^{(1)})$ of a tiling of finite local complexity is now constructed as follows: the level n vertices are the r_n-patches and its root represents the empty patch ($r_0 = 0$). There is a (vertical) edge between an r_n-patch and any of its extensions to an r_{n+1}-patch and all edges arise in this way.

The tree of patches is the Michon tree of the transversal of the tiling equipped with an ultrametric of the form

$$d(\xi, \eta) = \inf\{\delta(r_n) \,:\, \xi_n = \eta_n, \text{ with } r_n = |\xi_n|\}\,, \tag{7}$$

where $\delta : \mathbb{R}^+ \to \mathbb{R}^+$ is any strictly decreasing function converging to 0 at $+\infty$.

No particular structure of the tiling seems to point to a natural choice for the horizontal edges, the function δ above, or for determining the length of edges, or the choice function τ. These data have to be chosen according to the specific situation in order to define suitable approximating graph G_τ and δ.

Definition 3.4. By an *ordinary transverse spectral triple* for a tiling we mean a spectral triple for its canonical transversal (as in Def. 2.8) equipped with an ultrametric of the form (7).

The spectral triple depends on a choice of horizontal edges $\mathcal{H}$, a choice of strictly decreasing function $\delta : \mathbb{R}^+ \to \mathbb{R}^+$ tending to 0, and a choice function.

3.3. Substitution tilings

We consider now aperiodic primitive substitution tilings of finite local complexity. To these we may apply the construction of Section 3.2 to obtain an ordinary transverse spectral triple. But the extra structure coming from the substitution map allows one also to consider another spectral triple, namely the spectral triple of a one-sided subshift of finite type given by the substitution graph (see Section 3.1.1). The advantage of this latter approach lies in the fact that it can be extended into the longitudinal direction and therefore will provides us with a spectral triple for the continuous hull. We follow [40].

Consider a finite set $\mathcal{A} = \{t_i : i \in \mathcal{G}^{(0)}\}$ of translationally non-congruent tiles (called *prototiles*) in $\mathbb{R}^d$ indexed by a finite set $\mathcal{G}^{(0)}$. A *substitution* Φ on $\mathcal{A}$ with expansion factor $\theta > 1$ is a decomposition rule followed by an expansion by θ, namely Φ assigns to each prototile a patch of tiles in $\mathbb{R}^d$ with the properties: for each $i \in \mathcal{G}^{(0)}$, every tile in $\Phi(t_i)$ is a translate of an element of $\mathcal{A}$; and the subset of $\mathbb{R}^d$ covered by $\Phi(t_i)$ is the subset covered by t_i stretched by the factor θ. Such a substitution naturally extends to patches and even tilings whose elements are translates of the prototiles. It satisfies $\Phi(P - t) = \Phi(P) - \theta t$.

A patch P is *allowed* for Φ if there is an $m \geq 1$, an $i \in \{1, \ldots, k\}$, and a $v \in \mathbb{R}^N$, with $P \subset \Phi^m(t_i) - v$. The *substitution tiling space* associated with Φ is

the collection Ω_Φ of all tilings T of $\mathbb{R}^d$ such that every finite patch in T is allowed for Φ. Ω_Φ is not empty and, since translation preserves allowed patches, $\mathbb{R}^d$ acts on it by translation.

We assume that the substitution Φ is *primitive*, that is, for each pair $\{t_i, t_j\}$ of prototiles there is a $k \in \mathbb{N}$ so that a translate of t_i occurs in $\Phi^k(t_j)$. We also assume that all tilings of Ω_Φ are *non-periodic* and have FLC. It then follows that Ω_Φ is compact in the standard topology for tiling spaces and that $\Omega_\Phi = \Omega_T := \overline{\{T - x : x \in \mathbb{R}^d\}}$ for any $T \in \Omega_\Phi$. It also implies that the substitution map Φ seen as a map $\Omega_\Phi \to \Omega_\Phi$ is a homeomorphism (in particular it is bijective).

The substitution graph of Φ is the finite oriented graph $\mathcal{G} = (\mathcal{G}^{(0)}, \mathcal{G}^{(1)})$ whose vertices $\mathcal{G}^{(0)}$ stand for the indices of prototiles and whose edges encode the position of tiles in supertiles. More precisely, given two prototiles t_i and t_j the supertile $\Phi(t_i)$ may contain several tiles of type t_j (that is, tiles which are translationally congruent to t_j). These tiles are at different positions in the supertile $\Phi(t_i)$. For each possible position we introduce one oriented edge $\epsilon \in \mathcal{G}^{(1)}$ with range $r(\epsilon) = i$ and source $s(\epsilon) = j$.

The canonical transversal of a substitution tiling (alternatively a substitution subshift) can be encoded by a one-sided shift of finite type. To do so, one defines a map from the transversal Ξ to itself by desubstitution: if $T \in \Xi$, the origin is a puncture inside of a tile t_i of T. Therefore, in $\Phi^{-1}(T)$, the origin is inside of a unique tile t_j, whose puncture x_j doesn't need to be at the origin. Let $T' := \Phi^{-1}(T) - x_j \in \Xi$. The map $T \mapsto T'$ is a continuous surjection $\Xi \to \Xi$. For a given tiling T, the map $T \mapsto T'$ defines an edge in the substitution graph of Φ given by the position of t_i in the supertile $\Phi(t_j)$. By iteration, we get a continuous map $\Xi \to \Pi_\infty$ which commutes with the desubstitution map on Ξ and the shift on Π_∞ respectively. This map is called the Robinson map [25, 30] and denoted $\mathcal{R}$. Under an additional assumption on the substitution (called border-forcing), it is a homeomorphism.[6] Given t_i the tile at the origin of a tiling T and t_j the tile at the origin of $\Phi^{-n}(T)$, then t_i is a tile included in the patch $\Phi^n(t_j)$, or we can say that t_i is included in the n-supertile of type t_j. Remark how each finite path of length n corresponds to a tile at the origin (of type given by the source of the path) included in a n-supertile (given by the range of the path).

Example. The substitution graph of the Fibonacci substitution subshift of Example 3.1 is precisely the graph $\mathcal{G}$ giving rise to Figure 2 where vertices and edges have the following interpretation: The vertices correspond to the letters of the subshift. The arrows pointing towards the left vertex correspond to the occurrences of a and b in the substitution of a (which is baa), while those pointing towards the right vertex correspond to the occurrences of a and b in the substitution of b (which is ba) the dot showing which letter is concerned. The graph thus looks as

[6] Given a substitution tiling space Ω, it is always possible to find a substitution which produces this space and forces its border. It involves decorating the prototiles (hence increasing their number).

follows:

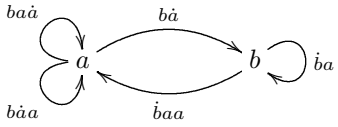

With the interpretation of paths on the substitution graph in terms of positions of tiles (here letters) in supertiles Figure 2 becomes the following tree of words:

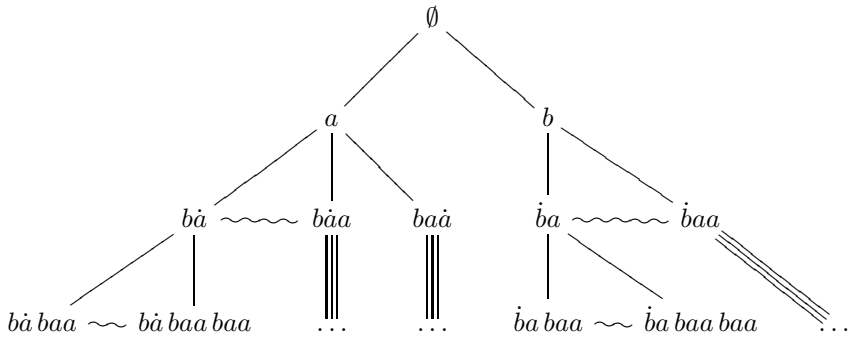

FIGURE 3. Part of the tree of words arising from the Fibonacci substitution. The dot indicates the letter (tile) in the nth order supertile which corresponds to the origin. For example $b\dot{a}\,baa$ stands for the supertile of order two $\Phi^2(b) = \Phi(b)\Phi(a)$ centered about the a which lies in the supertile of order one $\Phi(b) = ba$.

The same construction can be done for a tiling which is not in the transversal: if the origin belongs to a unique tile in $\Phi^n(T)$ for all $n \in \mathbb{Z}$, one can define similarly an element $\mathcal{R}(T) \in \Pi_{-\infty,\infty}(\mathcal{G})$, where $\Pi_{-\infty,\infty}(\mathcal{G})$ denotes the bi-infinite sequences over $\mathcal{G}$. It is of course not defined on all of Ω, but on a dense subset. It is injective, however, and its inverse can be extended to a continuous surjection $q : \Pi_{-\infty,\infty}(\mathcal{G}) \to \Omega$. The interpretation is the following: given $\gamma \in \Pi_{-\infty,\infty}$, the indices $(\gamma_i)_{i\geq 0}$ define a tiling $T \in \Xi$ by the inverse of the Robinson map. The indices $(\gamma_j)_{-N<j<0}$ define a Nth order microtile of type $s(\gamma_{-N})$ inside of the tile $s(\gamma_0)$. As N grows, it defines a decreasing sequence of microtiles which converge to a point x. Then $T - x$ is the tiling corresponding to γ. It can happen that two sequences of microtiles converge to the same point, meaning q is not injective.

We may suppose[7] that the substitution has a fixed point T^* such that the union over n of the nth order supertiles of T^* on 0 covers $\mathbb{R}^d$. Then $q^{-1}(\{T^*\})$ contains a single path and this path is constant, that is, the infinite repetition of a loop edge which we choose to be ϵ^*. And we pick a choice function $\hat{\tau}$ on $\mathcal{G}$, as in Section 3.1.1, which induces an embedding $\tau : \Pi(\mathcal{G}) \to \Pi_\infty(\mathcal{G})$ (if γ is a finite path, then $\tau(\gamma)$ is an infinite path which eventually is an infinite repetition of the edge ϵ^*).

[7]This can always be achieved by going over to a power of the substitution.

Let Ξ_t be the acceptance domain of prototile t (the set of all tilings in Ξ which have t at the origin). Note that the sets $\Xi_{t_v} \times t_v$ (for $v \in \mathcal{G}^{(0)}$) cover Ω. Their intersection turns out to have measure 0. Let $\Pi^v_{-\infty,\infty}$ be the set of bi-infinite paths which pass through v at level 0. Then $q_v : \Pi^v_{-\infty,\infty} \to \Xi_{t_v} \times t_v$ is a continuous almost one-to-one surjection.

In the section, we will use the chair substitution as an example. It is primitive and aperiodic but does not force its border. The Robinson map can still be defined, but is not a homeomorphism (*i.e.*, the one-sided shift of finite type is not an accurate representation of the transversal). For the sake of simplicity, we will nevertheless use it as an illustration.

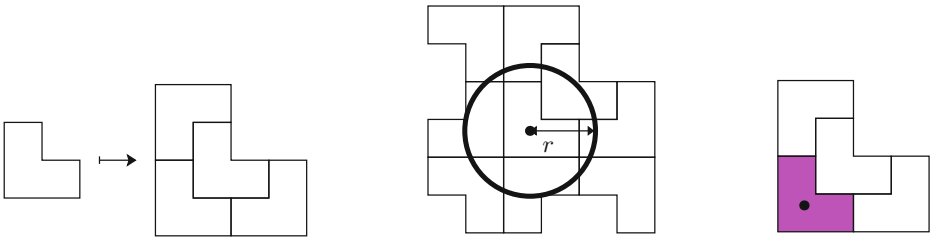

FIGURE 4. On the left, the chair substitution rule (with four prototiles, one for each orientation). In the middle, a vertex of the usual tree of patches associated with radius r. On the right, a vertex of the self-similar tree of patches as described in this section: a tile (colored) inside of a first-order supertile (equivalently, it is a length-one path in the substitution graph).

3.3.1. Transverse substitution spectral triple of a substitution tiling. We use now the positive part $q_v^+ : \Pi^v_{0,+\infty}(\mathcal{G}) \to \Xi_{t_v}$ of q_v to construct a second spectral triple for the transversal of a substitution tiling. Note that q_v^+ corresponds to the inverse of a restriction of $\mathcal{R}$ and thus is a homeomorphism.

We choose fundamental (transversal) horizontal edges

$$\hat{\mathcal{H}}_{\rm tr} \subset \left\{ (\epsilon, \epsilon') \in \mathcal{G}^{(1)} \times \mathcal{G}^{(1)} \,:\, \epsilon \neq \epsilon',\ s(\epsilon) = s(\epsilon') \right\}$$

which we suppose to satisfy the condition

(C) if $s(\epsilon) = s(\epsilon')$ there is a path of edges in $\hat{\mathcal{H}}_{\rm tr}$ linking ϵ with ϵ'.

This condition implies that the corresponding approximating graph will be connected.

What does an edge $(\gamma\epsilon, \gamma\epsilon') \in \mathcal{H}_{{\rm tr},n}$ stand for? Let's consider first the case $n = 0$ (so there is no γ). The two paths $\tau(\epsilon), \tau(\epsilon') \in \Pi_\infty(\mathcal{G})$ both start at vertex v and differ on their first edge. At some level $n_{(\epsilon,\epsilon')}$ they become equal again; let's say that $v_{(\epsilon,\epsilon')} \in \mathcal{G}^{(0)}$ is the vertex at which this happens. The edge (ϵ, ϵ') therefore defines a pair (η, η') of paths of length $n_{(\epsilon,\epsilon')}$ which have the same source and the

same range vertex $v_{(\epsilon,\epsilon')}$ and otherwise differ on each edge; notably η and η' are the first $n_{(\epsilon,\epsilon')}$ edges of $\tau(\epsilon)$ and $\tau(\epsilon')$, respectively. So the information encoded by (ϵ,ϵ') is the $n_{(\epsilon,\epsilon')}$-supertile corresponding to $v_{(\epsilon,\epsilon')}$ together with the position of two tiles of type t_v encoded by the paths η and η'. Let us denote by $r_{(\epsilon,\epsilon')} \in \mathbb{R}^d$ the vector of translation from the first to the second tile of type t_v. The remaining common part of the paths $\tau(\epsilon)$ and $\tau(\epsilon')$ (eventually an infinite repetition of ϵ^*) places the $n_{(\epsilon,\epsilon')}$-supertile corresponding to $v_{(\epsilon,\epsilon')}$ into some translate of T^*. And $r_{(\epsilon,\epsilon')}$ does not depend on this part.

Now the situation for $n > 0$ is similar, the only difference being that the paths $\tau(\gamma\epsilon)$ and $\tau(\gamma\epsilon')$ now split at the nth vertex and meet for the first time again at level $n + n_{(\epsilon,\epsilon')}$. If we denote by $r_{(\gamma\epsilon,\gamma\epsilon')} \in \mathbb{R}^d$ the translation vector between the encoded tiles then, due to self-similarity, one has:

$$r_{(\gamma\epsilon,\gamma\epsilon')} = \theta^{|\gamma|} r_{(\epsilon,\epsilon')}. \tag{8}$$

See Figure 5 for an illustration.

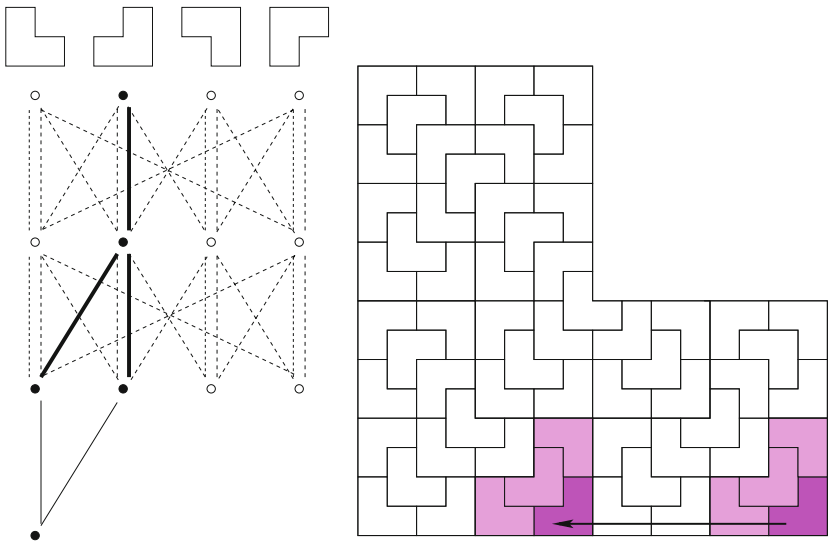

FIGURE 5. A translation defined by a horizontal arrow $h \in \mathcal{H}_{\mathrm{tr},3}$. The arrow represents the vector r_h. Here $n = 2$ (the paths have lengths 2), and $n_h = 1$ (the paths join further down at level $n + n_h = 3$).

We take $(\delta_n)_n$ of exponentially decreasing form: $\delta_n = \delta_{tr}^n$, where $\delta_{tr} \in (0,1)$ is a parameter which may be adapted. We moreover fix a choice function τ. Theorem 2.5 provides us with a self similar spectral triple $(C(\Xi_{t_v}), D_{\mathrm{tr}}^v, \mathfrak{H}_{\mathrm{tr}}^v)$ for the algebra $C(\Xi_{t_v})$. Since q_v^+ is a homeomorphism we could also use the version of the spectral triple of Theorem 2.4. We call the triple the *transverse substitution spectral triple for the prototile* t_v of the substitution tiling.

Definition 3.5. By a *transverse substitution spectral triple* of a substitution tiling we mean the direct sum over $v \in \mathcal{G}^{(0)}$ of the transverse spectral triples for the prototiles t_v defined as above.

The spectral triple depends on a choice of fundamental horizontal edges $\hat{\mathcal{H}}_{\rm tr}$ satisfying condition (C), a parameter δ_{tr} determining the length function, and a choice function. Since $\hat{\mathcal{H}}_{\rm tr}$ satisfies condition (C) above and $(\delta_n)_n$ is exponentially decreasing [39], the Connes distance induces the topology of Ξ_v. A transverse substitution spectral triple is thus a second spectral triple for the canonical transversal Ξ.

3.3.2. Longitudinal spectral triples for the prototiles of a substitution tiling. We now use the negative part of q_v, $q_v^- : \Pi^v_{-\infty,0}(\mathcal{G}) \to t_v$ to construct a spectral triple which we call longitudinal. Notice that $\Pi^v_{-\infty,0}$ can be identified with $\Pi^v_{\infty}(\widetilde{\mathcal{G}})$, where $\widetilde{\mathcal{G}}$ is the graph obtained from $\mathcal{G}$ by reversing the orientation of its edges: one simply reads paths backwards, so follows the edges along their opposite orientations. We choose a subset $\hat{\mathcal{H}}_{\rm lg} \subset \left\{(\widetilde{\varepsilon},\widetilde{\varepsilon}') \in \widetilde{\mathcal{G}^{(1)}} \times \widetilde{\mathcal{G}^{(1)}} \,:\, \widetilde{\varepsilon} \neq \widetilde{\varepsilon}',\ s(\widetilde{\varepsilon}) = s(\widetilde{\varepsilon}')\right\}$ again satisfying condition (C). To obtain the interpretation of a longitudinal horizontal edges it is more useful to work with reversed orientations, that is, view $\hat{\mathcal{H}}_{\rm lg} \subset \left\{(\varepsilon,\varepsilon') \in \mathcal{G}^{(1)} \times \mathcal{G}^{(1)} \,:\, \varepsilon \neq \varepsilon',\ r(\varepsilon) = r(\varepsilon')\right\}$, as this was the way the Robinson map $\mathcal{R}$ was defined. Then $(\varepsilon,\varepsilon')$ with $r(\varepsilon) = r(\varepsilon')$ determines a pair of microtiles (t,t') of type $s(\varepsilon)$ and $s(\varepsilon')$, respectively, in a tile of type $r(\varepsilon)$. The remaining part of the double path $(\widetilde{\tau}(\varepsilon),\widetilde{\tau}(\varepsilon'))$ serves to fix a point in the two microtiles. Of importance is now the vector of translation $a_{(\epsilon,\epsilon')}$ between the two points of the microtiles.

Similarily, an edge in $\mathcal{H}_{{\rm lg},n}$ will describe a pair of $(n+1)$th order microtiles in an nth order microtile. By self-similarity again, the corresponding translation vector $a_{(\gamma\epsilon,\gamma\epsilon')} \in \mathbb{R}^d$ between the two $(n+1)$th order microtiles will satisfy

$$a_{(\gamma\epsilon,\gamma\epsilon')} = \theta^{-|\gamma|} a_{(\epsilon,\epsilon')}\,. \tag{9}$$

See Figure 6 for an illustration.

Again we use an exponential decreasing length function $\delta_n = \delta_{lg}^n$, where δ_{lg} is another parameter. Moreover we set $\widetilde{\varepsilon}^* = \varepsilon^*$ and choose a choice function $\widetilde{\tau}$. q_v^- is however not injective and the topology of $\Pi^v_{-\infty,0}$ differs from that of t_v (the first is totally disconnected whereas the second is connected). It is therefore a priori not clear that Theorem 2.5 provides us with a spectral triple for the algebra $C(t_v)$. However, (with the notation $C(\partial\mathcal{T})_0$ introduced in eq. 3) it can be shown that ${q^*}^{-1}(C(\Pi^v_{-\infty,0})_0)$ contains all functions over t_v which are Hölder continuous with exponent $\alpha = \frac{-\log(\delta_{lg})}{\log(\theta)}$. Thus Theorem 2.5 provides us with a spectral triple $(C(t_v), D^v_{\rm lg}, \mathfrak{H}^v_{\rm lg})$ for the prototile t_v.

Definition 3.6. By a *longitudinal substitution spectral triple* for the prototile t_v we mean a spectral triple as defined above.

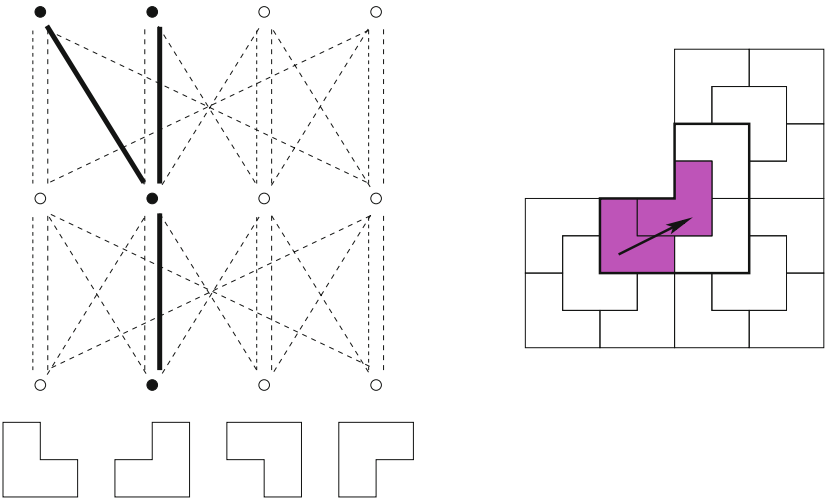

FIGURE 6. A microtile pattern associated with a horizontal arrow $h \in \mathcal{H}_{\mathrm{lg},2}$ (the pattern shown has the size of a single tile). The arrow represents the vector a_h.

The spectral triple depends on a choice of fundamental horizontal edges $\hat{\mathcal{H}}_{\mathrm{lg}}$ satisfying condition (C), a parameter δ_{lg} determining the length function, and a choice function. It should be noted that the Connes distance of this spectral triple does not induce the topology of t_v.

3.3.3. The full substitution spectral triple of a substitution tiling. We now combine the above triples to get a spectral triple $(C(\Omega_\Phi), \mathfrak{H}, D)$ for the whole tiling space Ω_Φ. The graphs $\mathcal{G}$ and $\widetilde{\mathcal{G}}$ have the same set of vertices $\mathcal{G}^{(0)}$, so we notice that the identification

$$\Pi_{-\infty,+\infty}(\mathcal{G}) = \bigcup_{v\in\mathcal{G}^{(0)}} \Pi^v_{-\infty,0}(\mathcal{G}) \times \Pi^v_{0,+\infty}(\mathcal{G}) = \bigcup_{v\in\mathcal{G}^{(0)}} \Pi^v_{\infty}(\widetilde{\mathcal{G}}) \times \Pi^v_{\infty}(\mathcal{G})$$

suggests to construct the triple for Ω_Φ as follows: first we can use the tensor product construction for spectral triples to obtain a spectral triple for $C(t_v \times \Xi_{t_v}) \cong C(t_v) \otimes C(\Xi_{t_v})$ from the two spectral triples considered above. Furthermore, the C^*-algebra $C(\Omega_\Phi)$ is a subalgebra of $\bigoplus_{v\in\mathcal{G}^{(0)}} C(t_v \times \Xi_{t_v})$ and so the direct sum of the tensor product spectral triples for the different tiles t_v provides us with a spectral triple for $C(\Omega_\Phi)$. Its Hilbert space, representation and Dirac operator are given by

$$\mathfrak{H} = \bigoplus_{v\in\mathcal{G}^{(0)}} \mathfrak{H}^v_{\mathrm{tr}} \otimes \mathfrak{H}^v_{\mathrm{lg}}\,, \quad \pi = \bigoplus_{v\in\mathcal{G}^{(0)}} \pi^v_{\mathrm{tr}} \otimes \pi^v_{\mathrm{lg}}\,, \quad D = \bigoplus_{v\in\mathcal{G}^{(0)}} \left(D^v_{\mathrm{tr}} \otimes \mathbf{1} + \chi \otimes D^v_{\mathrm{lg}}\right), \tag{10}$$

where χ is the grading of the transversal triple (which comes from flipping the orientations of the edges in $\hat{\mathcal{H}}_{\mathrm{tr}}$).

Definition 3.7. By a *full substitution spectral triple* of a substitution tiling we mean a spectral triple for Ω_Φ as defined above.

The spectral triple depends on a choice of transverse and longitudinal fundamental horizontal edges $\hat{\mathcal{H}}_{\rm tr}, \hat{\mathcal{H}}_{\rm lg}$ satisfying condition (C), two parameters δ_{tr}, δ_{lg} determining the length function, and a choice function.

4. Zeta functions and spectral measures

The Dirac D operator of a spectral triple is supposed to have compact resolvent. As we suppose for simplicity that 0 is not in its spectrum, the sequence of eigenvalues of $|D|^{-1}$, ordered decreasingly and counted with their multiplicity, tends to zero at infinity. In the construction of Section 2.2.2 the eigenvalues are given by the length function so that the zeta function for the spectral triple of Section 2.2.2 is formally given by the series

$$\zeta(z) = \sum_{(v,v')\in\mathcal{H}} \delta(v,v')^z\,, \tag{11}$$

or even by

$$\zeta(z) = \sum_{n=1}^{\infty} \#\mathcal{H}_n\ \delta_n^z \tag{12}$$

in case when the edges can be written $\mathcal{H} = \bigcup_n \mathcal{H}_n$ such that the length function depends only on the level n of the edge, *i.e.*, $\delta(v,v') = \delta_n$ for all $v, v' \in \mathcal{H}_n$. In the latter case ζ is manifestly independent of the choice function (however notice that in Palmer's construction the length function depends on the choice function and hence so does ζ).

We suppose that the series has a finite abscissa of convergence $s_0 > 0$, i.e., is convergent for z with $\Re(z) > s_0$. This so-called metric dimension s_0 has a certain significance for the examples we discussed. There is a general expectation that a good spectral triple for a metric space has a metric dimension which coincides with a dimension of the space. Such a result holds for the spectral triple of Section 2.3 for a compact ultrametric space when the minimal choice for the horizontal edges is employed. This is a case in which the length function depends only on the level and so the zeta function is independent of the choice function.

Theorem 4.1 ([50]). *Let (X,d) be a compact ultrametric space. If its associated Michon tree has uniformly bounded branching*[8]*, then the Pearson–Bellissard spectral triple has metric dimension equal to the upper box dimension of X.*

Similarly, the spectral measure, defined on functions $f \in C(X)$ by

$$\mathcal{T}(f) = \lim_{s\to s_0^+} \frac{1}{\zeta(s)} \mathrm{Tr}(|D|^{-s}\pi(f))$$

[8]The number of edges in $e \in \mathcal{T}^{(1)}$ with source vertex $s(e) = v \in \mathcal{T}^{(0)}$ is uniformly bounded in v.

is not just any measure but yields a measure well-known to the cases. We have already mentionned Palmer's theorem (Theorem 2.6) expressing the fact that for a compact metric space one can construct a spectral triple for which s_0 is the Hausdorff dimension, and the spectral measure is the Hausdorff measure.

4.1. The metric dimension and complexities

In the context of ordinary transverse spectral triples defined for tilings (based on the tree of patches), or of spectral triples defined for subshifts (based on the tree of words), the metric dimension is related to complexities.

Recall that we call r-patch a patch of the tiling which has a puncture at the origin, and just covers $B_r(0)$, that is, all tiles of the patch intersect $B_r(0)$ non-trivially. As we assume that our tiling is FLC we can define the function $p : \mathbb{R}^+ \to \mathbb{N}$ associating to r the number of r-patches. It is called the *patch counting* (or *complexity*) function. This function might increase exponentially but we are here interested in tilings where it grows only polynomially, a feature which can be interpreted as a sign of order in the tiling. We call

$$\underline{\beta} = \sup\{\gamma : p(r) \geq r^\gamma \text{ for large } r\}\,, \quad \overline{\beta} = \inf\{\gamma : p(r) \leq r^\gamma \text{ for large } r\} \tag{13}$$

the lower and upper complexity exponents. Notice that these exponents can be alternatively defined as the lim inf and lim sup of $\log p(r)/\log r$ respectively. In case both are equal we call their common value the *weak complexity exponent*. If $p(r)$ is equivalent to cr^β for some constant $c > 0$, then β is simply called the complexity exponent. This implies that $\underline{\beta} = \overline{\beta} = \beta$ but is stronger than existence of the weak exponent.

One could imagine different definitions of r-patches using other geometric objects than balls or other norms on $\mathbb{R}^d$ than the Euclidean one. This would effect the function $p(r)$ but not the exponents.

The weak complexity exponents have an important interpretation. If one takes $\delta(r) = \frac{1}{r}$ in the definition for the ultra metric of the transversal Ξ of an FLC tiling, then the weak complexity exponents are the lower and upper box dimensions of Ξ [27]:

$$\underline{\beta} = \underline{\dim}(\Xi)\,, \qquad \overline{\beta} = \overline{\dim}(\Xi)\,.$$

In the case of (primitive) substitution tilings of $\mathbb{R}^d$, the complexity exponent exists [28] and is equal to the Hausdorff dimension of Ξ [29], we comment on that again further down.

4.1.1. The case of ordinary transverse spectral triples for tilings. Let $a(v)$ be the *branching number* of the vertex v **minus one**, so $a(v)+1$ is the number of edges in $\mathcal{T}^{(1)}$ which have source v. We define the kth related Dirichlet series

$$\zeta_k(s) = \sum_{v \in \mathcal{T}^{(0)}} a(v)^k\, \delta(r_v)^s\,, k \in \mathbb{N}\,,$$

where $r(v)$ is the radius of the patch associated with the vertex $v \in \mathcal{T}^{(0)}$, and we use the convention $0^0 = 0$. Then the zeta functions of the ordinary transverse spectral triples constructed from $\mathcal{H}$ in Section 3.2 are given as follows:

- $\zeta^{\max} := (\zeta_1 + \zeta_2)/2$ is the zeta function if we take the maximal choice $\mathcal{H} = \mathcal{H}^{\max}$;
- $\zeta^{\min} := \zeta_0$ is the zeta function for a minimal choice $\mathcal{H} = \mathcal{H}^{\min}$.

We denote by $s_0^{\min\,/\,\max}$ the abscissa of convergence of $\zeta^{\min\,/\,\max}$.

Theorem 4.2 ([39]). *Consider the ordinary transverse spectral triple of a d-dimensional tiling of finite local complexity* (Def. 3.4)*. Assume that the function δ belongs to $L^{1+\epsilon}([0,\infty)) \setminus L^{1-\epsilon}([0,\infty))$ for all ϵ small enough. Then*

$$\underline{\beta} \leq s_0^{\min} \leq s_0^{\max} \leq \overline{\beta} + d - 1\,.$$

The theorem applies for instance to spectral triples constructed using $\delta(r) = \frac{1}{r+1}$ as function. The latter does not belong to $L^1([0,\infty))$ but to $L^{1+\epsilon}([0,\infty))$ for all $\epsilon > 0$.

4.1.2. The case of spectral triples defined for subshifts. The complexity function for an infinite word is $p(n)$ = number of distinct factors of length n. We define further the *right-special complexity* $p_{\rm rs}$ as:

$$p_{\rm rs}(n) = \text{number of distinct right-special factors of length } n\,, \tag{14}$$

as well as the *privileged complexity* $p_{\rm pr}$ as:

$$p_{\rm pr}(n) = \text{number of distinct privileged factors of length } n\,. \tag{15}$$

One defines weak complexity exponents $\beta_{\rm pr}$ and $\beta_{\rm rs}$ for $p_{\rm pr}$ and $p_{\rm rs}$ just as in equation (13) for p. Then we have (Corollary 6.2 in [33])

$$\beta_{\rm pr} \leq \beta_{\rm rs} = \beta - 1\,.$$

When do we have equality between the exponents $\beta_{\rm pr} = \beta_{\rm rs}$? This question is related to the notion of *almost finite rank*. Given an r-patch q of an FLC tiling T, the Delone set $\mathcal{L}_q$ of occurrences of q in T can be tiled by Voronoi cells with finitely many prototiles. Let $n(q)$ be this number. If $n(q)$ is bounded in q, then the tiling T is said to have *finite rank*. If there are constants $a, b > 0$ such that $n(q) \leq a \log(r(q))^b$, where $r(q)$ is the size of the patch, then T is said to have *almost finite rank*.

We denote by $\zeta^{\rm pr}$ the zeta function associated with the privileged choice of horizontal edges.

Theorem 4.3 ([33]). *Consider the spectral triple of a one-sided subshift as in Def.* 3.2*. If the subshift is repetitive, has almost finite rank, and $\delta \in L^{1+\epsilon}([0,\infty)) \setminus L^{1-\epsilon}([0,\infty))$ for all ϵ small enough, then $\zeta^{\max}$ and $\zeta^{\rm pr}$, as well as the related Dirichlet series ζ_k, have the same abscissa of convergence. If furthermore the subshift admits weak complexity exponents, then $\beta_{\rm pr} = \beta_{\rm rs} = \beta - 1$, and β coincides with the common abscissa of convergence.*

As an aside we mention that the equality $\beta_{\rm pr} = \beta_{\rm rs}$ can be seen as an asymptotic version of a stronger result which holds for rich words [21]: $p_{\rm pr}(n)+p_{\rm pr}(n+1) = p(n+1) - p(n) + 2$. Indeed one has $p_{\rm pr}(n) \leq p(n+1) - p(n) \leq (|\mathcal{A}| - 1)p_{\rm pr}(n)$,

and thus $\beta_{\rm pr} = \beta_{\rm rs}$. And for rich words, privileged factors and palindromes are equivalent.

Example. For the Fibonacci subshift of Example 3.1, the tree of words $\mathcal{T}$ (shown at the end of Section 3.1) has a single branching vertex in $v_n^* \in \mathcal{T}_n^{(0)}$ for each n, whose branching number is exactly 2 (this is a property shared by all Sturmian words). That is $a(v_n^*) = 1$ for all n, and $a(v) = 0$ for all other $v \neq v_n^*$ for all n. We choose the weight δ so that it only depends on the length of the factors: for instance $\delta(v) = \frac{1}{|v|}$ to satisfy the hypothesis of Theorem 4.3. Then for all $k \in \mathbb{N}$ the zeta functions are all equal to

$$\zeta_k(z) = \sum_{n\geq 1} \sum_{v\in\mathcal{T}_n^{(0)}} a(v)^k \delta(v)^z = \sum_{n\geq 1} \delta(v_n^*)^z = \sum_{n\geq 1} \frac{1}{n^z},$$

whose abscissa of convergence is 1. This is also the complexity exponent of the Fibonacci subshift (all Sturmian words have indeed complexity function $p(n) = n + 1$).

4.2. The zeta function for a self-similar spectral triple

We determine in more detail the form of the zeta-function for the triple associated with a subhift of finite type with self-similar choices of horizontal edges, length function and choice function. The most important of these choices is the scale $\rho \in (0, 1)$, i.e., the parameter such that $\delta_n = \rho^n$. Let A be the *graph matrix* of $\mathcal{G}$, that is, the matrix A with coefficients A_{vw} equal to the number of edges which have source v and range w. The number of paths of length n starting from v and ending in w is then $A^n{}_{vw}$. We require that A is *primitive*: $\exists N \in \mathbb{N}, \forall v, w, A_{vw}^N > 0$. Under this assumption, A has a non-degenerate positive eigenvalue $\Lambda_{\rm PF}$ which is strictly larger than the modulus of any other eigenvalue. This is the Perron–Frobenius eigenvalue of A. Let $\lambda_1, \lambda_2, \ldots, \lambda_p$ be the eigenvalues of A, ordered by decreasing modulus, and with $\lambda_1 = \Lambda_{\rm PF}$. We now compute the zeta function $\zeta(z)$ as in equation (12). The cardinality of $\#\mathcal{H}_n$ can be estimated as

$$\#\mathcal{H}_n = \sum_{v,w\in\mathcal{G}^{(0)}} A_{vw}^{n-1} n_w = \sum_{j=1}^{p} C_{\hat{\mathcal{H}}}^j \lambda_j^n + o(|\lambda_p|)\,,$$

where $n_w = \#\{(\epsilon, \epsilon') \in \hat{\mathcal{H}} : s(\epsilon) = s(\epsilon') = w\}$, and the $C_{\hat{\mathcal{H}}}^j$, $1 \leq j \leq p$, are constants. Hence each eigenvalue of A gives a geometric contribution to the Dirichlet series $\zeta(z)$.

Example. For the Fibonacci subshift of Example 3.1, the graph matrix reads $A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, with eigenvalues $\lambda_1 = \tau^2$ (the Perron–Frobenius), and $\lambda_2 = \tau^{-2}$, where $\tau = (\sqrt{5} + 1)/2$ is the golden mean. We choose $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{\max}$, so that

$$\#\mathcal{H}_n = \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}^{n-1} \begin{pmatrix} 3 \\ 2 \end{pmatrix} = (2 + \tau)\,\tau^{2n} + 5\tau^{-2n},$$

(we counted only unoriented edges here) and we have exactly

$$\zeta(z) = \sum_{n\geq 1} \#\mathcal{H}_n \ \rho^{nz} = \frac{2+\tau}{1-\tau^2\rho^z} + \frac{5}{1-\tau^{-2}\rho^z}.$$

So the spectral dimension of the self-similar spectral triple for the Fibonacci tiling is $s_0 = \frac{2\log\tau}{-\log\rho}$.

In general we can show the following:

Theorem 4.4. *Consider a self-similar spectral triple as in Definition* 3.3 *with scale* ρ*. Suppose that the graph matrix is diagonalizable with eigenvalues* λ_j*,* $j = 1, \dots, p$*. The zeta-function* ζ *extends to a meromorphic function on* $\mathbb{C}$ *which is invariant under the translation* $z \mapsto z + \frac{2\pi\imath}{\log\rho}$*. It is given by*

$$\zeta(z) = \sum_{j=1}^{p} \frac{C^j_{\hat{\mathcal{H}}}}{1-\lambda_j\rho^z} + h(z)$$

where h *is an entire function. In particular* ζ *has only simple poles which are located at* $\{\frac{\log\lambda_j + 2\pi\imath k}{-\log\rho} : k \in \mathbb{Z}, j = 1, \dots, p\}$ *with residues given by*

$$\mathrm{Res}(\zeta, \frac{\log\lambda_j + 2\pi\imath k}{-\log\rho}) = \frac{C^j_{\hat{\mathcal{H}}}\lambda_j}{-\log\rho}\,. \tag{16}$$

In particular, the metric dimension is equal to $s_0 = \frac{\log\Lambda_{\mathrm{PF}}}{-\log\rho}$*.*

Remark 4.5. The periodicity of the zeta function with purely imaginary period whose length is only determined by the factor ρ comes from the self-similarity of the construction. It is a feature which distinguishes our spectral triples from known triples for manifolds. Note also that ζ may have a (simple) pole at 0, namely if 1 is an eigenvalue of the graph matrix A.

Remark 4.6. The location of the poles and hence also the metric dimension does not depend on the choice of $\hat{\mathcal{H}}$ (neither on the choice function).

Remark 4.7. In the general case, when A is not diagonalizable, it is no longer true that the zeta-function has only simple poles. Indeed, at least if $|\lambda_i| \neq 1$ the zeta function $\zeta(z)$ has poles of order m_j at $z = \frac{\log\lambda_j + 2\pi\imath k}{-\log\rho}$ where m_j is the size of the largest Jordan block of A corresponding to eigenvalue λ_j.

Theorem 4.4 applies to the transversal and the longitudinal spectral triples discussed in Section 3.3.1 and 3.3.2. In particular it applies to the spectral triple used in [28]. It is the transverse substitution spectral triple of Section 3.3.1 with minimal choice for the horizontal edges and $\rho_{tr} = \theta^{-1}$. This shows that the transverse metric dimension for a substitution tiling of $\mathbb{R}^d$ equals d.

The metric dimension is additive under the tensor product construction of spectral triples. Hence the metric dimension of the spectral triples for Ω_Φ from Section 3.3.3 is

$$s_0 = d\Big(\frac{\log\theta}{-\log\rho_{tr}} + \frac{\log\theta}{-\log\rho_{lg}}\Big).$$

5. Laplacians

Recall that a spectral triple $(C(X), D, \mathfrak{H})$ for a compact space X together with a state $\mathcal{T}$ on the Hilbert space $\mathfrak{H}$ defines a quadratic form

$$(f, g) = \mathcal{T}\big([D, \pi(f)]^*[D, \pi(g)]\big)$$

which may be interpreted as the analog of a Laplacian on X. This is of particular interest if X does not carry an a priori differentiable structure.

But the Laplacian does not come for free. In fact, subtle analytic questions have to be resolved which depend on further choices. The question is whether the quadratic form can be extended to a quadratic form on the Hilbert space $L^2_{\mathbb{R}}(X, \mu)$, where μ is typically (but not necessarily) the spectral measure. This problem involves the determination of a core for the form (a dense subspace on which it can be defined, and then extended by a closure operation). The resulting form will a priori depend on the choice for the core. One then has to check that the closure of the form has the desired properties of a Dirichlet form.

Although the usual procedure of construction of such forms employs the spectral state $\mathcal{T}$ we will proceed at first in a slightly different manner. Let tr be a trace on the Hilbert space $\mathfrak{H}$ and consider the bi-linear form

$$(f, g) \mapsto \mathrm{tr}\big(|D|^{-s}[D, \pi(f)]^*[D, \pi(g)]\big) \tag{17}$$

for $f, g \in C(X)$ such that $|D|^{-s}[D, \pi(f)]^*[D, \pi(g)]$ is trace class. We consider two cases here. In both cases the spectral triple is defined by an approximating graph $G_\tau = (V, E)$ as in Section 2.2.2.

(1) We consider in Section 5.1 $\mathrm{tr} = \mathrm{Tr}$ *the usual operator trace on* $\mathfrak{H}$ and leave s as a parameter. Upon averaging over choice functions τ we obtain the Pearson–Bellissard Dirichlet form and its Laplacian. We discuss its spectral theory in the case that X is an arbitrary compact ultrametric space getting more concrete results for $X = \Xi$, the canonical transversal of a tiling space or a subshift.

(2) In Section 5.2 we employ $\mathrm{tr} = \mathrm{str}$ *the singular trace on* $\mathfrak{H}$ at $s = s_0$. This corresponds to using the spectral state $\mathcal{T}$ for the definition of the quadratic form. We get very concrete results in the case that $X = \Omega_\Phi$ is a continuous Pisot substitution tiling space, obtain two forms (and associated Laplacians): one of longitudinal nature, and one of transversal nature. The two forms may be combined into a single one.

5.1. The operator trace and the Pearson–Bellissard Laplacian

In this section we let (X, d) be a compact ultrametric space. We refer the reader to Section 2.3 where the canonical Michon tree $\mathcal{T}$ is constructed.

We wish to extend the quadratic form (17) using the operator trace $\mathrm{tr} = \mathrm{Tr}$. The naive idea, namely to consider the graph Laplacian on the approximating graph $G_\tau = (V, E)$ cannot work, as it is defined on $\ell^2(V)$ and no continuous non trivial continuous function on X restricts to $\ell^2(V)$. The way around this problem is to average over choice functions τ. For this we need a probability measure on

the set of choice functions. Let $\tau : \mathcal{T}^{(0)} \to \partial\mathcal{T}$ be a choice function. We can either say that τ chooses an infinite extension for every finite path, or we can say that it chooses the follow up vertex for every vertex. As such, τ can be thought of as an infinite family of finite choices, namely for each vertex which one to choose next. The family of choice functions can therefore be identified with the product $Y = \prod_{v\in\mathcal{T}^{(0)}} \mathcal{T}^{(0)}(v)$, where, we recall, $\mathcal{T}^{(0)}(v)$ is the finite set of vertices following v (hence can be ignored if v is not branching). Furthermore, there is a one-to-one correspondence between Borel probability measures on X and product probability measures on Y ([39] Lemma 5.8). We endow Y with a product probability measure $\mathbb{P}$, and let ν be the corresponding measure on X. We divide equation (17) by 2 (to avoid counting unoriented edges twice) and average over choices to define the quadratic form on $L^2_{\mathbb{R}}(X,\nu)$:

$$Q_s(f,g) = \frac{1}{2}\int_Y \mathrm{Tr}\big(|D|^{-s}[D,\pi_\tau(f)]^*[D,\pi_\tau(g)]\big)\, d\mathbb{P}(\tau)\,. \tag{18}$$

It follows from Pearson–Bellissard's work that equation (18) defines a Dirichlet form whose domain is generated by (real-valued) locally constant functions on X of the form $\chi_v = 1\!\!1_{q([v])}$, where $[v] \subset \partial\mathcal{T}$ is the set of all infinite paths going through v. And one identifies $Q_s(f,g)$ with $\langle f, \Delta_s g\rangle_{L^2_{\mathbb{R}}(X,\nu)}$ for the *Laplacian* Δ_s: a non-positive definite self-adjoint operator on $L^2_{\mathbb{R}}(X,\nu)$, with pure point spectrum.

Let us comment on the form of equations (17) and (18). Choice functions can be interpreted as analogue to tangent vectors over the "noncommutative manifold" X, and hence Y stands for the unit tangent sphere bundle to X. The noncommutative gradient of a function $\nabla_\tau f = [D, \pi_\tau(f)]$ therefore stands for the directional derivative of f along τ. The quadratic form (18) is therefore reminiscent to the classical integral $\int_M \mathrm{g}(\nabla f, \nabla g)\, d\mathrm{vol} = \langle f, -\Delta g\rangle_{L^2(M)}$ defining the Laplace–Beltrami operator over a Riemannian manifold (M, g). The analogy is not perfect, however, as integration of a function over the manifold, $\int_M f\, d\mathrm{vol}$, usually corresponds to the application of the spectral state.

We now particularise the above construction to determine explicitly Δ_s. We consider the *minimal choice* of horizontal edges, $\mathcal{H} = \mathcal{H}^{\min}$, as described in Section 2.3. We assume further that μ is the spectral measure, and that the length function satisfies: $\delta(h) = d(q\circ\tau(s(h)), q\circ\tau(r(h)))$, *i.e.*, the distance in X between the source and range of the image of $h \in \mathcal{H}$ in X. Then this yields the spectral triple of Pearson–Bellissard as in Definition 2.8. If all branching vertices of the tree $\mathcal{T}$ have exactly two outgoing vertical edges, then our general construction reduces to that of Pearson–Bellissard in [50]. For instance, this is the case for X a Sturmian subshift as in Example 3.1, with $\mathcal{T}$ its tree of words.

We now average the form in equation (18) over minimal choices of edges $\mathcal{H}$, to obtain the Pearson–Bellissard Laplacian. It can be determined explicitly and diagonalised [28]. There are two regimes: if $s < s_0 + 2$ then Δ_s is unbounded, while for $s > s_0 + 2$ it is bounded. (In the bounded case it is possible to obtain an embedding of the transversal of substitution tiling space in a Euclidean space [29].) The eigenvalues and eigenvectors are parametrized by the set $\mathcal{T}^{(0)}_{br} \subset \mathcal{T}^{(0)}$ of

branching vertices: a basis of eigenvectors of Δ_s is given by functions

$$\varphi_v = \frac{1}{\mu[u]}\chi_u - \frac{1}{\mu[u']}\chi_{u'}, \quad v \in \mathcal{T}^{(0)}_{br}, \quad u \neq u' \in \mathcal{T}^{(0)}(v),$$

where $\chi_v = \mathbb{1}_{q([v])}$ and $\mu[v]$ stands for the μ-measure of $q([v])$. The associated eigenvalues λ_v can be calculated explicitly. For $s = s_0$, the spectral dimension of X, one can compute the Weyl asymptotics of the eigenvalues of Δ_{s_0}: the number of eigenvalues of modulus less than λ behaves asymptotically like $\lambda^{s_0/2}$, in analogy with the classical case.

Example. For the Fibonacci subshift of Example 3.1, the eigenfunction, associated with the vertex *aba* in its tree $\mathcal{T}$ shown at the end of Section 3.1, reads

$$\varphi_{aba} = \frac{1}{\mu[abab]}\chi_{abab} - \frac{1}{\mu[abaa]}\chi_{abaa}.$$

We now turn to the self-similar case of substitution tilings: $X = \Xi_\Phi$ is the discrete transversal to a substitution tiling space with substitution map Φ. Let A be the substitution matrix of the substitution Φ, $\mathcal{G}$ its substitution graph, and $\mathcal{T}$ the associated tree (of super tiles), as in Sections 3.3 and 3.1.1. The spectral measure μ of the corresponding transverse spectral triple (see Section 3.3.1) is the transverse invariant ergodic probability measure on Ξ_Φ.

The Laplacian Δ_s is in this case completely and explicitly determined [28]: all eigenvectors and corresponding eigenvalues $(\varphi_v, \lambda_v), v \in \mathcal{T}^{(0)}_{br}$, can be computed algorithmically from the ones (φ_0, λ_0) associated with the root of $\mathcal{T}$, by use of the Cuntz–Krieger algebra $\mathcal{O}_A$ of the substitution matrix. The Cuntz–Krieger algebra $\mathcal{O}_A$ [15] is the universal C*-algebra generated by partial isometries $U_1, U_2, \ldots, U_n$ of an infinite-dimensional separable Hilbert space as follows: the operators $U_1, U_2, \ldots, U_n$ are in one-to-one correspondence with the vertices $u_1, \ldots, u_n$ in $\mathcal{T}^{(0)}_1$ (*i.e.*, the prototiles), and subject to the relations $U_i^* U_i = \sum_j A_{ij} U_j U_j^*$. There are two faithful $*$-representations of $\mathcal{O}_A$: $\rho_1 : \mathcal{O}_A \to L^2(X, \mu)$ and $\rho_2 : \mathcal{O}_A \to \ell^2(\mathcal{T}^{(0)})$ such that

$$\varphi_v = \rho_1(U_{i_m} \cdots U_{i_2} U_{i_1})\varphi_0\,, \qquad \lambda_v = \langle \rho_2(U_{i_m} \cdots U_{i_2} U_{i_1})\delta_0, \delta_v \rangle\, \lambda_0\,,$$

where we wrote $v = u_{i_m}, \ldots, u_{i_2}, u_{i_1}$ the unique sequence of vertices in $\mathcal{T}^{(0)}$ from v up to the root, the root vertex not being included. Moreover, δ_0, δ_v are the basis elements of $\ell^2(\mathcal{T}^{(0)})$ associated with the root and vertex v and $\langle \cdot, \cdot \rangle$ denotes the scalar product in $\ell^2(\mathcal{T}^{(0)})$.

5.2. The singular trace and Laplacians for substitution tilings

We now consider the quadratic form in (17) which we obtain if we use for tr a singular trace str on $\mathfrak{H}$ at parameter $s = s_0$. A priori there may be many singular traces, but if the limit below exists then, apart from an overall factor, they all

amount to:

$$\begin{aligned} Q(f,g) &= \mathrm{str}\big(|D|^{-s_0}[D,\pi(f)]^*[D,\pi(g)]\big) \\ &= \lim_{s\to s_0^+} \frac{1}{\zeta(s)} \mathrm{Tr}\big(|D|^{-s}[D,\pi(f)]^*[D,\pi(g)]\big)\,. \end{aligned} \tag{19}$$

In other words $Q(f,g) = \mathcal{T}\big([D,\pi(f)]^*[D,\pi(g)]\big)$ is defined by the spectral state $\mathcal{T}$.

Our main application is to a self similar spectral triple from a substitution tiling, as discussed in Section 3.3.3. Note that in (19) we do not integrate over choice functions. Such an integration would in fact not change much, as there are only finitely many possible choices due to our self-similarity constraint on the choice function.

Recall that $C(\Omega_\Phi)$ is a subalgebra of $\oplus_{t\in\mathcal{A}} C(t\times\Xi_t)$ where $\mathcal{A}$ is the set of prototiles. We can employ the tensor product structure of $C(t\times\Xi_t)\cong C(t)\otimes C(\Xi_t)$ to decompose the form into the sum of a transversal form and a longitudinal form:

$$Q(f,g) = Q_{lg}(f,g) + Q_{tr}(f,g)\,, \tag{20}$$

where

$$Q_\alpha(f,g) = \mathcal{T}_\alpha\big([D_\alpha,\pi_\alpha(f)]^*[D_\alpha,\pi_\alpha(g)]\big)\,. \tag{21}$$

Here $\alpha = tr$ or lg and the corresponding objects $\mathcal{T}_\alpha, D_\alpha, \pi_\alpha$ are the spectral state, the Dirac operator and the representation associated with the transverse and longitudinal substitution spectral triple of t of Sections 3.3.1 and 3.3.2, respectively. There are subtle technicalities behind this decomposition, in particular for the spectral state (which holds only for so-called strongly regular operators on $\mathfrak{H}$ [40]) which we will not discuss here.

In both cases the operator $[D_\alpha,\pi_\alpha(f)]^*[D_\alpha,\pi_\alpha(g)]$ is diagonal and gives a contribution $(\delta_e^\alpha f)^*\delta_e g$ on the edge $e\in E_n^\alpha = q\circ\tau\times q\circ\tau(\mathcal{H}_{\alpha,n})$, where

$$\delta_e^\alpha f = \frac{f(r(e)) - f(s(e))}{\rho_\alpha^n}\,. \tag{22}$$

Moreover

$$Q_\alpha(f,g) = \lim_{s\to s_0^+} \frac{1}{\zeta_\alpha(s)} \sum_{n\geq 1} \#E_n^\alpha \rho_\alpha^{ns}\; q_n^\alpha(f,g)$$

with

$$q_n^\alpha(f,g) = \frac{1}{\#E_n^\alpha} \sum\nolimits_{e\in E_n^\alpha} \overline{\delta_e^\alpha f}\;\delta_e^\alpha g\,.$$

Notice that $\lim_{s\to s_\alpha^+} \frac{1}{\zeta_\alpha(s)} \sum_{n\geq 1} \#E_n^\alpha \rho_\alpha^{ns} = 1$, and hence we have $Q_\alpha(f,g) = \lim_n q_n^\alpha(f,g)$ provided the limit exists. We now briefly explain how we can evaluate these limits. The following two paragraphs are a bit technical, the reader will find the main result stated and discussed in the last paragraph.

The longitudinal form. Given a fundamental edge $(\epsilon, \epsilon') \in \hat{\mathcal{H}}_{lg}$, recall that $a_{(\epsilon,\epsilon')} \in \mathbb{R}^d$ denotes the translation vector between the punctures of the microtiles associated with ϵ and ϵ'. If $a_e \in \mathbb{R}^d$ denotes the corresponding vector for $e \in E_{lg,n}$ of type (ϵ, ϵ'), that is, $e = q \circ \tau \times q \circ \tau(\gamma\epsilon, \gamma\epsilon')$ for some γ of length n, then by (9) we have $a_e = \theta^{-n} a_h$, so $s(e) = r(e) + \theta^{-n} a_h$. We make for large n the Taylor approximation in equation (22)

$$\delta_e^{lg} f \simeq \left(\frac{\theta^{-1}}{\rho_{lg}}\right)^n \; (a_h \cdot \nabla) f(s(e))\,.$$

If $\rho_{lg} = \theta^{-1}$, and we further approximate the sum of $(\delta_e^{lg} f)^* \delta_e^{lg} g$ over $e \in E_n^{lg}$ by a Riemann integral, then for n large $q_n^{lg}(f, g)$ gives contributions of the form $\int_t (a_h \cdot \nabla)\bar{f}\; (a_h \cdot \nabla) g \; d\mu_{lg}^t$.

We define the operator on $L^2_{\mathbb{R}}(\Omega_\Phi, d\mu)$:

$$\Delta_{lg} = c_{lg} \nabla_{lg}^\dagger \mathcal{K} \nabla_{lg}\,, \quad \text{with} \quad \mathcal{K} = \sum_{t \in \mathcal{A},\, h \in \mathcal{H}_1^{lg}(t)} \text{freq}(t)\; a_h \otimes a_h\,, \tag{23}$$

where $\hat{\mathcal{H}}_{lg}(t) = \{(\epsilon, \epsilon') : s(\epsilon) = s(\epsilon') = t\}$, freq$(t)$ is the frequency of tile t, and we write ∇_{lg} for the *longitudinal gradient* on Ω_Φ. The latter takes derivatives along the leaves of the foliation: $\nabla_{lg} = \mathbf{1} \otimes \nabla_{\mathbb{R}^d}$. This leads to the expression

$$Q_{lg}(f, f) = \begin{cases} \langle f,\; \Delta_{lg}\; f \rangle_{L^2_{\mathbb{R}}(\Omega_\Phi, d\mu)} & \text{if } \rho_{lg} = \theta^{-1} \\ 0 & \text{if } \rho_{lg} > \theta^{-1} \\ +\infty & \text{if } \rho_{lg} < \theta^{-1} \end{cases}, \quad \text{for all } f \in C^2_{lg}(\Omega_\Phi)\,, \tag{24}$$

where $C^2_{lg}(\Omega_\Phi)$ is the space of longitudinally C^2 functions on Ω_Φ. So we see that for $\rho_{lg} \geq \theta^{-1}$, Δ_{lg} is essentially self-adjoint on the domain $C^2_{lg}(\Omega_\Phi)$, and therefore the form Q_{lg} is closable. For $\rho_{lg} < \theta^{-1}$ the form is not closable.

The transversal form. Given a fundamental edge $h = (\epsilon, \epsilon') \in \hat{\mathcal{H}}_{tr}$ let $t = s(\epsilon) = s(\epsilon')$ and denote as before by $r_h \in \mathbb{R}^d$ the return vector between the occurrences of t in the supertiles associated with ϵ and ϵ'. If $r_e \in \mathbb{R}^d$ denotes the corresponding vector for $e \in E_n^{tr}$ of type (ϵ, ϵ'), then by self-similarity we have $r_e = \theta^n r_h$, so $s(e) = r(e) + \theta^n r_h$.

We assume now that the substitution is Pisot, *i.e.*, θ is a Pisot number. Then we know that there are plenty of dynamical eigenfunctions: continuous functions f_β satisfying $f_\beta(\omega + r) = e^{2\imath\pi\beta(r)} f_\beta(\omega)$ for some $\beta \in \mathbb{R}^{d^*}$ and all $\omega \in \Omega_\Phi$ and $r \in \mathbb{R}^d$. In fact, the set of $\beta \in \mathbb{R}^{d^*}$ for which such a function exists is a dense subgroup of $\mathbb{R}^{d^*}$ and the Pisot substitution conjecture[9] states that if the substitution matrix is irreducible then the dynamical spectrum is purely discrete which means that the eigenfunctions generate all of $L^2(\Omega_\Phi, \mu)$. We assume this to be the case, and we choose the linear span of dynamical eigenfunctions to be the core of Q_{tr}. We have

$$\delta_e^{tr} f_\beta = \frac{1}{\rho_{tr}^n}\big(f_\beta(s(e) + \theta^n r_h) - f_\beta(s(e))\big) = \frac{1}{\rho_{tr}^n}\left(e^{2\imath\pi\theta^n\beta(r_h)} - 1\right) f(s(e))\,.$$

[9] discussed at length in Chapter 2 of this book

By the arithmetic properties of Pisot numbers $\theta^n\beta(r_h)$ tends to an integer as n goes to infinity. Moreover, the speed of convergence is governed by the Galois conjugates of θ of greatest modulus: $\theta_j, j = 2, \ldots, L$, $|\theta_j| = |\theta_2|$. It follows that, for n large, we have

$$\delta_e^{tr} f_\beta \simeq \left(\frac{\theta_2}{\rho_{tr}}\right)^n \left(\sum\nolimits_{j=2}^{L} p_{\beta(r_h)}(\theta_j)\right)\ f(s(e))\,,$$

where $p_{\beta(r_h)}$ is some polynomial with rational coefficients. We are left with summing the terms $(\delta_e^{tr} f)^* \delta_e^{tr} g$ over $e \in E_n^{tr}$, and approximate by a Riemann sum to get an asymptotically equivalent expression for $q_n^{tr}(f, g)$. There are averaging subtleties coming from the phases α_j of the θ_j and we have to assume that[10]: $\alpha_j - \alpha_{j'} + 2k\pi + 2\pi k' \frac{\log \rho_{tr}}{\log \rho_{lg}} \neq 0$, $\forall k, k' \in \mathbb{Z}$. Define the operator Δ_{tr} on the linear space of dynamical eigenfunctions by

$$\Delta_{tr} f_\beta = -c_{tr}(2\pi)^2 \sum_{t\in\mathcal{A}, h\in\hat{\mathcal{H}}_{tr}(t)} \mathrm{freq}(t) \sum_{j=2}^{L} |p_{\beta(h)}(\theta_j)|^2\ f_\beta\,. \tag{25}$$

Then, on the space of dynamical eigenfunctions the transversal form is given by

$$Q_{tr}(f_\beta, f_\beta) = \begin{cases} \langle f_\beta, \Delta_{tr} f_\beta\rangle_{L^2(\Omega_\Phi, d\mu)} & \text{if } \rho_{tr} = |\theta_2| \\ 0 & \text{if } \rho_{tr} > |\theta_2| \\ +\infty & \text{if } \rho_{tr} < |\theta_2| \end{cases}.$$

Clearly, Q_{tr} is closable but trivial if $\rho_{tr} > |\theta_2|$, whereas Q_{tr} is not closable if $\rho_{tr} < |\theta_2|$.

Main result and geometric interpretation. We summarize here the results about the Dirichlet forms. For a Pisot number θ of degree $J > 1$, we denote $\theta_j, j = 2, \ldots, J$, the other Galois conjugates in decreasing order of modulus. We write the subleading eigenvalues in the form $\theta_j = |\theta_2| e^{\imath\alpha_j}$, $2 \leq j \leq L$, where $\alpha_j \in [0, 2\pi)$. In particular, $|\theta_j| < |\theta_2|$ for $j > L$.

Theorem 5.1 ([40]). *Consider a Pisot substitution tiling of $\mathbb{R}^d$ with Pisot number θ of degree $J > 1$. Assume that for all $j \neq j' \leq L$ one has*

$$\alpha_j - \alpha_{j'} + 2\pi k + 2\pi \frac{\log|\theta_2|}{\log\theta} k' \neq 0\,, \quad \forall k, k' \in \mathbb{Z}\,. \tag{26}$$

Set $\rho_{lg} = \theta^{-1}$ and $\rho_{tr} = |\theta_2|$.

If the dynamical spectrum is purely discrete then the set of finite linear combinations of dynamical eigenfunctions is a core for Q on which it is closable. Furthermore, $Q = Q_{tr} + Q_{lg}$, and $Q_{tr/lg}$ has generator $\Delta_{tr/lg} = \sum_{h\in\mathcal{H}_{tr/lg,1}} \Delta_{tr/lg}^h$ given by

$$\Delta_{lg}^h f_\beta = -c_{lg}(2\pi)^2 \mathrm{freq}(t_h)\beta(a_h)^2 f_\beta,$$
$$\Delta_{tr}^h f_\beta = -c_{tr}(2\pi)^2 \mathrm{freq}(t_h)\langle \widetilde{r_h}^\star, \beta\rangle^2 f_\beta$$

[10]The ratio $\log(\theta)/\log(|\theta_2|)$ is irrational unless θ is a unimodular Pisot number of degree $J = 3$ [56].

where t_h is the tile associated with (the source of the vertical edges linked by) h, and the constants c_{lg} and c_{tr} depend only on the substitution matrix.

We now explain the term $\widetilde{r_h}^\star$ in the above equation, and give an interpretation of the Laplacians $\Delta_{tr/lg}$ as elliptic second-order differential operators with constant coefficients on the maximal equicontinuous factor of the dynamical system $(\Omega_\Phi, \mathbb{R}^d)$.

We assume for simplicity that θ is unimodular, see [40] for the general case. Our assumption that the dynamical spectrum is purely discrete is known to be equivalent to the tiling being a cut-and-project tiling. The maximal equicontinuous factor of the dynamical system $(\Omega_\Phi, \mathbb{R}^d)$ coincides with the dJ-torus $\mathbb{T}$ of the torus parametrisation of the cut-and-project scheme. The substitution induces a hyperbolic homeomorphism on that torus which allows us to split the tangent space at each point into a stable and an unstable subspace, S and U. The unstable tangent space is d-dimensional and can be identified with the space in which the tiling lives. We write $\widetilde{r}$ for the vector in U corresponding to r via the identification of U with the space in which the tiling lives. The stable space S can be split further into eigenspaces of the hyperbolic map, namely $S = S_2 + S'$ where S_2 is the sum of eigenspaces of the Galois conjugates of θ which are next to leading in modulus (θ_j for $j = 2, \ldots, L$). Finally $^\star : U \to S_2 \subset S$ is the reduced star map. This is Moody's star map followed by a projection onto S_2 along S'.

Since the dynamical spectrum is pure point and all eigenfunctions continuous the factor map $\pi : \Omega_\Phi \to \mathbb{T}$ induces an isomorphism between $L^2(\Omega, \mu)$ and $L^2(\mathbb{T}, \eta)$, where η is the normalized Haar measure on $\mathbb{T}$. The Dirichlet form Q can therefore also be regarded as a form on $L^2(\mathbb{T}, \eta)$. Now the directional derivative at x along $u \in U \oplus S$ is given by

$$(\langle u, \nabla \rangle f_\beta)(x) = \frac{d}{dt} f_\beta(x + tu) \,|_{t=0} \, = 2\pi i \langle u, \beta \rangle f_\beta(x).$$

And we thus have

$$\Delta_{lg}^h = c_{lg} \mathrm{freq}(t_h) \langle \widetilde{a}_h, \nabla \rangle^2,$$
$$\Delta_{tr}^h = c_{tr} \mathrm{freq}(t_h) \langle \widetilde{r_h}^*, \nabla \rangle^2.$$

To summarize, the Dirichlet form Q viewed on $L^2(\mathbb{T}, \eta)$ has as generator the Laplacian $\Delta = \sum_{h \in \hat{\mathcal{H}}_{lg}} \Delta_{lg}^h + \sum_{h \in \hat{\mathcal{H}}_{tr}} \Delta_{tr}^h$ on $\mathbb{T}$ which is a second-order differential operator with constant coefficients containing (second) derivatives only in the directions $U + S_2$.

6. Characterization of order

In this section we explain how noncommutative geometry can be used to characterise combinatorial properties of tilings and subshifts. These properties, equidistribution of frequencies and bounded powers, are signs of aperiodic order. In this theory, which has been proposed in [39] and extended in [33], the interest is no

longer Rieffel's question after the construction of a spectral triple whose associated Connes distance induces the topology, but a comparison of the different distance functions which arise for different choices of the choice functions.

6.1. Notions of aperiodic order

After recalling some notions of aperiodic order we focus on the spectral triples of Sections 3.1 and 3.2 with the aim to derive criteria for high aperiodic order (Section 6.2). We will define these notions mainly for tilings of $\mathbb{R}^d$. For symbolic tilings (infinite words) the corresponding notions are easily adapted.

Complexities and complexity exponents. We introduced the complexity function $p(r)$ and various complexity exponents in Sections 4.1 and 4.1.2. The smaller the growth of the complexity function, the more ordered the tiling appears. An ordered tiling is expected to have a subexponential or polynomially bounded complexity function.

Repetitiveness. The repetitiveness function $R : \mathbb{R}^+ \to \mathbb{R}^+$ for a tiling T is defined as follows: $R(r)$ is the smallest r' such that any r'-patch of T contains an occurrence of all the r-patches of T.

We will assume that $R(r)$ is finite for all $r \geq 0$. In the context of tilings of finite local complexity this means that the tiling is *repetitive* and is equivalent to the minimality of the dynamical system $(\Omega, \mathbb{R}^d)$. A well-ordered tiling is also supposed to have a low repetitiveness function. The repetitiveness function is related to the complexity function in several ways. For instance, Lagarias and Pleasants established a bound $R(r) \geq cp(r)^{1/d}$, for some $c > 0$ and all r large, which holds in general for any repetitive tiling of $\mathbb{R}^d$ [41]. Among repetitive tilings, the *linearly repetitive* ones [20, 17] are commonly regarded as the most ordered [41]. These are tilings for which there exists a constant $c_{\mathrm{LR}} > 0$ such that

$$R(r) \leq c_{\mathrm{LR}}\, r\,, \quad \forall r \geq 0\,.$$

For linearly repetitive tilings, the reverse inequality to Lagarias–Pleasants' bound holds too [43], and one has: $c\, p(r)^{1/d} \leq R(r) \leq c'\, p(r)^{1/d}$ for some $c, c' > 0$ and r large.

Finite and almost-finite ranks, defined in Section 4.1, are related notions. They characterize bounded or slow increasing number of repetition types.

Repulsiveness and bounded powers. A return vector to a patch in a tiling is any vector joining (the punctures of) two occurrences of that patch in the tiling. A repetitive tiling is said to be *repulsive* if any return to an r-patch grows at least like r, for r large. Linear repetitive tilings are repulsive [43, 9], and repulsiveness is also a signature of high aperiodic order. For example, a non-repulsive tiling has patches of arbitrarily large sizes which overlap on arbitrarily large parts. These overlaps force local periodicity, as is easily seen in dimension 1 and for words.

An infinite word with language $\mathcal{L}$ is repulsive if its index of repulsiveness,

$$\ell = \inf\Big\{ \frac{|W| - |w|}{|w|} \,:\, w, W \in \mathcal{L},\, w \text{ is a proper prefix and suffix of } W \Big\}\,, \qquad (27)$$

is positive: $\ell > 0$. This is equivalent to *bounded powers*: there exists an integer p such that each factor occurs at most p times consecutively: $\forall u \in \mathcal{L}, u^{p+1} \notin \mathcal{L}$.

Equidistributed frequencies. A uniquely ergodic tiling is said to have *equidistributed frequencies* if the frequency of any r-patch behaves like a given function of r, for r large. As this function is independent of the patches, it is easily related to the complexity. Namely, a tiling has equidistributed frequencies if there are constants $c, c' > 0$ such that for any r-patch P one has:

$$c\, p(r)^{-1} \leq \mathrm{freq}(P) \leq c'\, p(r)^{-1}\,. \tag{28}$$

A linearly repetitive tiling has equidistributed frequencies ([39] Theorem 1.8).

Uniform bound on the number of patch extensions. In a tiling with finite local complexity, any r-patch P has finitely many extensions: if $r = r_n$, there are finitely many r_{n+1}-patches containing P. If this number is uniformly bounded in n, then the tiling is said to have a *uniform bound on the number of patch extensions*. Tilings with FLC in dimension 1, and words over a finite alphabet, obviously have such a uniform bound. This notion is closely related to finite rank, defined in Section 4.1.

A tiling with equidistributed frequencies and for which the complexity function satisfies $p(4r) \leq c\, p(r)$, for some $c > 0$ and all r large, has a uniform bound on the number of patch extensions ([39] Lemma 4.15).

6.2. A noncommutative geometrical criterion for aperiodic order

In the following we work with spectral triples for one-sided subshifts (Def. 3.2) and with ordinary transversal spectral triples for tilings or subshifts (Def. 3.4). Recall that these spectral triples depend on various choices, namely those for the horizontal edges and the length function, but also that for the choice function. While we keep the choices for the horizontal edges and the length function fixed we treat that for the choice function τ as a parameter and thus obtain a Connes distance $d_C = d_\tau$ which depends on this parameter. How are these distance functions related? In particular, considering the infimum and supremum of d_τ over all choices

$$d_{\sup} = \sup_\tau d_\tau\,, \qquad d_{\inf} = \inf_\tau d_\tau\,,$$

we may ask, are $d_{\sup}$ and $d_{\inf}$ equivalent in the sense of Lipschitz? By this we mean that there exists a constant $C > 0$ such that

$$C^{-1} d_{\inf} \leq d_{\sup} \leq C d_{\inf}.$$

This is the non commutative geometrical criterion which will characterize a certain form of aperiodic order, provided clever choices for the horizontal edges and the length function have been made.

We first consider repetitive aperiodic tilings of finite local complexity and their ordinary transverse spectral triples. We fix the *maximal* choice $\mathcal{H} = \mathcal{H}^{\max}$ for the horizontal edges. Identifying the canonical transversal Ξ of the tiling with

the set of infinite paths on the tree of patches we obtain the following expressions ([39] Section 4.1)

$$d_{\inf}(\xi,\eta) = \delta(|\xi\wedge\eta|)\,, \quad d_{\sup}(\xi,\eta) = \delta(|\xi\wedge\eta|) + \sum_{n>|\xi\wedge\eta|} \big(b_n(\xi)+b_n(\eta)\big)\,\delta(n)\,. \tag{29}$$

Here $\xi\wedge\eta$ is the greatest common prefix of the paths ξ and η and hence $|\xi\wedge\eta|$ is the radius of the largest r-patch the tilings corresponding to ξ and η have in common (around the origin). Furthermore $b_n(\xi)=1$ if the nth vertex of the path ξ is a branching vertex, otherwise $b_n(\xi)=0$.

The first formula says that $d_{\inf}$ corresponds to the metric defined by δ in (7). We need to assume that this function satisfies the following multiplicative inequalities, which are satisfied for instance for power functions:

$$\delta(ab) \le \bar{c}\,\delta(a)\delta(b) \tag{30a}$$

$$\delta(2a) \ge \underline{c}\,\delta(a) \tag{30b}$$

for some constants $\bar{c},\underline{c}>0$ and for all a, b large. The following result shows that Lipschitz equivalence of $d_{\inf}$ and $d_{\sup}$ is a necessary criterion for equidistribution of frequencies.

Theorem 6.1 ([39, Thm. 4.16]). *Consider a tiling which has equidistributed patch frequencies and whose complexity function satisfies: $p(4r)\le cp(r)$ for some $c>0$ and all r large. Suppose that δ satisfies the inequalities in equation* (30a)*. Then $d_{\inf}$ and $d_{\sup}$ are Lipschitz-equivalent.*

We now focus on subshifts considering the spectral triples for one-sided subshifts (Def. 3.2) with *privileged* horizontal edges $\mathcal{H}=\mathcal{H}^{\rm pr}$. In other words, $u_1,u_2\in\mathcal{T}^{(0)}$ are linked with a horizontal edge if and only if u_1 and u_2 are privileged words, and there is a privileged word $v\in\mathcal{T}^{(0)}$ such that u_1 and u_2 are two distinct complete first returns to v (see Section 3.1). As above the spectral triple requires a choice for the function δ, and will depend parametrically on the choice functions. We consider again the infimum $d_{\inf}$ and supremum $d_{\sup}$ of the spectral metric $d_C=d_\tau$ over all choice functions τ. These metrics, if continuous, are given on the one-sided subshift space Ξ by

$$d_{\inf}(\xi,\eta) = \delta(\|\xi\,\widetilde{\wedge}\,\eta\|)\,, \quad d_{\sup}(\xi,\eta) = \delta(\|\xi\,\widetilde{\wedge}\,\eta\|) + \sum_{n>\|\xi\wedge\eta\|} \big(\delta(\|\widetilde{\xi}_n\|)+\delta(\|\widetilde{\eta}_n\|)\big)\,. \tag{31}$$

Now $\xi\,\widetilde{\wedge}\,\eta$ denotes the greatest common *privileged* prefix of ξ and η. Furthermore we write $\|u\|=n$ if u is an nth-order privileged word, that is, obtained as nth iterated complete first return to the empty word. Finally $\widetilde{\xi}_n$ denotes the privileged prefix of nth order of ξ. In particular, $\|\xi\wedge\eta\|=m$ if $\widetilde{\xi}_m=\widetilde{\eta}_m$ but $\widetilde{\xi}_{m+1}\neq\widetilde{\eta}_{m+1}$.

Now Lipschitz equivalence of $d_{\inf}$ and $d_{\sup}$ is a necessary and sufficient criterion for bounded powers.

Theorem 6.2 ([33, Thm 5.1]). *Consider a minimal aperiodic one-sided subshift over a finite alphabet. Suppose that the function δ used to construct the above spectral triples satisfies the two inequalities* (30)*. Then the subshift has bounded powers if and only if $d_{\inf}$ and $d_{\sup}$ are Lipschitz-equivalent. Furthermore, the corresponding two-sided subshift has bounded powers if and only if the one-sided subshift has bounded powers.*

To emphasize the importance of the above result we discuss the special case of Sturmian subshifts (an example of which, the Fibonacci subshift, is given in Example 3.1). For these subshifts, having bounded powers is equivalent to linear repetitiveness. Sturmian subshifts depend on a parameter θ, the slope, whose number theoretical properties are reflected in various properties of the subshift. The following result says that our noncommutative geometric criterion characterizes properties of irrational numbers.

Corollary 6.3. *Consider a Sturmian subshift with slope θ. Suppose that the function δ used to construct the above spectral triples satisfies the two inequalities* (30)*. The following are equivalent:*

(i) *The subshift is linearly repetitive;*
(ii) *The continued fraction expansion of its slope θ has bounded coefficients;*
(iii) *The metrics $d_{\inf}$ and $d_{\sup}$ are Lipschitz-equivalent.*

One can say even a little more: If the coefficients in the continued fraction expansion of the slope grow very fast (see [39, Thm. 4.14] for a precise statement) then $d_{\sup}$ is not even continuous on the subshift space, that is, it is not compatible with its topology.

7. K-homology

We now discuss an application of the spectral triples for compact ultrametric spaces to the noncommutative topology of such spaces. According to the general theory a spectral triple over a C^*-algebra A gives rise to a K-homology class of that algebra which in turn defines a group homomorphism from the K-group of A to the group of integers [12]. A compact ultrametric space X is totally disconnected and its K_0-group $K_0(C(X))$ is isomorphic to $C(X,\mathbb{Z})$, and thus as a $\mathbb{Z}$-module generated by the indicator functions on the clopen sets of X. The K_1-group $K_1(C(X))$ is trivial. In this section we give an answer to the question, which group homomorphisms $C(X,\mathbb{Z}) \to \mathbb{Z}$ can be obtained from the spectral triples of a compact ultrametric set X.

7.1. Fredholm modules and their pairing with K-theory

A spectral triple is sometimes referred to as an *unbounded* Fredholm module. This suggests that it can somehow be reduced to a (bounded) Fredholm module.

Definition 7.1. An *odd* Fredholm module $(A, F, \mathfrak{H})$ over a C^*-algebra A is given by a Hilbert space $\mathfrak{H}$, a representation $\pi : A \to B(\mathfrak{H})$, and a bounded, self-adjoint operator F, such that $F^2 = 1$ and for all $a \in A$, $[F, \pi(a)]$ is compact.

An *even* Fredholm module is an odd Fredholm module, with an additional self-adjoint operator Γ which satisfies $\Gamma^2 = 1$, which commutes with each operator $\pi(a)$ and anti-commutes with F.

We see that the main difference to the definition of a spectral triple is that F is bounded and $F^2 = 1$. So it is no surprise that a spectral triple gives rise to a Fredholm module by applying a polar decomposition to D and taking F to be its unitary part[11], $D = F|D|$. Indeed, since D is self-adjoint we can obtain F from D by replacing the eigenvalues of D by their sign, $F = \mathrm{sign}(D)$.

As for even spectral triples, the grading operator of a Fredholm module allows one to decompose $\mathfrak{H} = \mathfrak{H}^+ \oplus \mathfrak{H}^-$, such that the representation acts diagonally as $\pi_+ \oplus \pi_-$, the operator Γ acts as $1 \oplus (-1)$, and

$$F = \begin{pmatrix} 0 & T^* \\ T & 0 \end{pmatrix}$$

for some unitary operator $T : \mathfrak{H}^+ \to \mathfrak{H}^-$.

Fredholm modules for A are at the basis of the construction of the K-homology group of A. We will not explain this construction here. Our interest lies in the group homomorphisms into $\mathbb{Z}$ the K-homology classes define on K-theory and these can be expressed directly on the level of Fredholm modules. Only even Fredholm modules lead to non-trivial homomorphisms on $K_0(A)$ and so for our purposes it is enough to consider even Fredholm modules.

Definition 7.2. Let F be a bounded operator between two Hilbert spaces. It is called a *Fredholm operator* if the dimension of its kernel and of its cokernel are finite. In this case, its index $\mathrm{ind}(F)$ is defined as

$$\mathrm{ind}(F) = \dim(\ker F) - \dim(\mathrm{coker}\, F).$$

The K_0-group $K(A)$ of a unital C^*-algebra is constructed from homotopy classes of projections p in A and in $M_n(A)$, the algebra of $n \times n$ matrices with entries in A (any natural n).

Theorem 7.3 ([12]). *Let $M = (A, F, \mathfrak{H})$ be an even Fredholm module and p be a projection in A. Then $\pi_-(p)T\pi_+(p)$ is a Fredholm operator between the two Hilbert spaces $\pi_+(p)\mathfrak{H}^+$ and $\pi_-(p)\mathfrak{H}^-$ and*

$$\varphi_M(p) = \mathrm{ind}\left(\pi_-(p)\, T\, \pi_+(p)\right) \tag{32}$$

induces a group homomorphism $\varphi_M : K_0(A) \to \mathbb{Z}$.

[11] We suppose for simplicity (and this is actually satisfied for our spectral triples) that D has no kernel.

We emphasize that the index has to be taken for the operator defined between the specified source and range spaces, and not all of $\mathfrak{H}^+$ and all of $\mathfrak{H}^-$.

Claiming that (32) induces a group homomorphism on the level of $K_0(A)$ means not only that it is additive, but also that $\varphi_M(p)$ is invariant under homotopy of projections, and that it extends (in a natural way) to projections in $M_n(A)$. We remark also that there is a natural notion of direct sum for Fredholm modules and that $\varphi_{M\oplus M'} = \varphi_M + \varphi_{M'}$. The above theorem may therefore be formulated also in the following way: The map $(M, p) \mapsto \varphi_M(p) \in \mathbb{Z}$ extends to a bi-additive map between the K-homology and the K-theory of A. This bi-additive map is called the *Connes pairing*.

7.2. Compact ultrametric space

We now consider the case of a compact ultrametric space X. We have seen in Section 2.3 that X can be identified with the space of paths on its Michon tree $\mathcal{T} = (\mathcal{T}^{(0)}, \mathcal{T}^{(1)})$ and obtain spectral triples depending on the choice of horizontal edges $\mathcal{H}$ and choice functions $\tau : \mathcal{T}^{(0)} \to \partial\mathcal{T}$. We fix an orientation $\mathcal{H} = \mathcal{H}^+ \cup \mathcal{H}^-$ which induces the grading on the Hilbert space $\mathfrak{H} = \mathfrak{H}^+ \oplus \mathfrak{H}^- = \ell^2(\mathcal{H}^+) \oplus \ell^2(\mathcal{H}^-)$. Theorem 2.4 now gives a spectral triple for $C(X) \cong C(\partial\mathcal{T})$.

Recall from Section 2.3 that the Dirac operator D of the spectral triple depends on the ultrametric. But note that $\mathrm{sign}(D)$ is independent of the that metric and hence the Fredholm module defined by the spectral triple depends only on $\mathcal{H}$ and τ. Indeed, $F = \mathrm{sign}(D)$ is given by

$$F = \begin{pmatrix} 0 & T^* \\ T & 0 \end{pmatrix},$$

where $T : \mathfrak{H}^+ \to \mathfrak{H}^-$ is induced by changing the orientation of the horizontal edge, $h \mapsto h^{\mathrm{op}}$. More precisely, if we denote $1_h \in \mathfrak{H}$ the function which is 1 on the edge $h \in \mathcal{H}$ then $T1_h = 1_{h^{\mathrm{op}}}$. We denote the associated homomorphism of the Fredholm module by $\varphi_{\tau,\mathcal{H}} : K_0(C(X)) \to \mathbb{Z}$, that is:

$$\varphi_{\tau,\mathcal{H}} : C(X, \mathbb{Z}) \longrightarrow \mathbb{Z}.$$

By compactness of X, any function in $C(X; \mathbb{Z})$ is a finite sum of the form $\sum \alpha_v \chi_v$, where χ_v is the characteristic function of the set of all infinite paths which pass through v. In particular, the K-theory is spanned by projections in $C(X)$ rather than in a matrix algebra. Therefore, $\varphi_{\tau,\mathcal{H}}$ is entirely determined by the values of $\{\varphi_{\tau,\mathcal{H}}(\chi_v)\ ;\ v \in \mathcal{V}\}$. These values can be explicitly computed using equation (32). Let $(1_h)_{h\in\mathcal{H}^+}$ be a basis for $\mathfrak{H}^+$. Then a simple computation shows that given $h \in \mathcal{H}^+$, $\pi_-(\chi_v)T\pi_+(\chi_v)\cdot 1_h = \chi_v(\tau \circ r(h))\chi_v(\tau \circ s(h))\cdot 1_{h^{\mathrm{op}}}$. Therefore, the index of this linear map (between the indicated Hilbert spaces) is:

$$\begin{aligned} \varphi_{\tau,\mathcal{H}}(\chi_v) = {} & \#\{h \in \mathcal{H}^+\ ;\ \chi_v(\tau \circ s(h)) \neq 0 \ \text{ and } \ \chi_v(\tau \circ r(h)) = 0\} \\ & -\#\{h \in \mathcal{H}^+\ ;\ \chi_v(\tau \circ r(h)) \neq 0 \ \text{ and } \ \chi_v(\tau \circ s(h)) = 0\}. \end{aligned} \tag{33}$$

The same arguments as in the proof for Lemma 2.3 show that whenever $h \in \mathcal{H}^+$ does not satisfy that either $r(h)$ or $s(h)$ lies on the unique path from the root to v,

then h does not belong to any of the sets occuring in equation (33). This equation can therefore be rewritten:

$$\varphi_{\tau,\mathcal{H}}(\chi_v) = \sum_{h\in\mathcal{H}^+} \Big(\chi_v(\tau\circ s(h)) - \chi_v(\tau\circ r(h))\Big), \tag{34}$$

the general term of the sum being non-zero only for finitely many h.

In particular, it results from these formulas that for any τ and any choice of minimal edges $\mathcal{H} = \mathcal{H}^{\min}$ we must have $\varphi_{\tau,\mathcal{H}}(1) = 0$, and that there are two vertices $v_1, v_2 \in \mathcal{T}_1^{(0)}$, such that $\varphi_\tau(\chi_{v_1}) = 1$ and $\varphi_\tau(\chi_{v_2}) = -1$. This provides the following proposition.

Proposition 7.4. *For any choice function τ and any choice of minimal edges $\mathcal{H}^{\min}$, the homomorphism $\varphi_{\tau,\mathcal{H}}$ is non-trivial. In particular, the K-homology class of a Pearson-Bellissard spectral triple is never trivial.*

Now, the question is: which homomorphisms on $K_0(C(X))$ can we obtain from Fredholm modules? The following proposition says that, if we only slightly relax our earlier assumptions by assuming that $\mathcal{H}$ can consist of several edges between the same vertices, then the condition found above, namely $\varphi(1) = 1$, is the only obstruction for an element of $\mathrm{Hom}(K_0(C(X)); \mathbb{Z})$ to come from a Fredholm module.

Proposition 7.5. *For any $\varphi \in \mathrm{Hom}(C(X;\mathbb{Z});\mathbb{Z})$ such that $\varphi(1) = 0$, there is a set of horizontal edges $\mathcal{H}$ (possibly bigger than what was defined as a maximal set), and a choice function τ such that $\varphi_{\tau,\mathcal{H}} = \varphi$.*

Proof. We begin by a simple remark. Suppose that we have two modules defined by two distinct choices $\tau, \mathcal{H}$ and $\tau', \mathcal{H}'$. If the horizontal edges coincide up to level n and the choice functions coincide up to level $n-1$ then $\varphi_{\tau,\mathcal{H}}(\chi_v) = \varphi_{\tau',\mathcal{H}_2}(\chi')$ for all vertices v of level smaller or equal to n. We can therefore construct inductively $\tau, \mathcal{H}$ from the values of $\varphi(\chi_v)$ to obtain $\varphi_{\tau,\mathcal{H}} = \varphi$ as follows.

Let $v_0 \in \mathcal{T}_0^{(0)}$ be the root. Choose $\tau(v_0)$ arbitrarily. Let $v_1, \ldots, v_k \in \mathcal{T}_1^{(0)}$. Suppose that v_1 is the vertex through which $\tau(v_0)$ passes. Then we must have $\tau(v_1) = \tau(v_0)$. For $i \neq 0, 1$ we may choose $\tau(v_i)$ arbitrarily. We now choose the horizontal edges $\mathcal{H}_1$ of level 1, cf. Figure 7. Let $n_1, \ldots, n_k$ be the numbers $n_i := \varphi(\chi_{v_i})$. Since $\sum_i \chi_{v_i} = 1$ and $\varphi(1) = 0$, the n_i sum to 0. By convention, we say that "n horizontal edges from v_i to v_j" consists of $|n|$ edges with source v_i and range v_j if n is positive, and $|n|$ edges with range v_i and source v_j if it is negative. Then $\mathcal{H}_1^+$ shall consist of n_1 edges from v_1 to v_2, $n_1 + n_2$ edges from v_2 to v_3, $\ldots$, and $\sum_{i=1}^{k-1} n_i$ edges from v_{k-1} to v_k. A simple computation shows that however we extend τ and whatever the choice of horizontal edges of higher level, $\varphi_{\tau,\mathcal{H}}(\chi_v)$ assumes the correct values.

To proceed with the inductive construction, assume that the choice function and horizontal edges are determined up to level $n-1$ and n, respectively. Let $w_0 \in \mathcal{T}_n^{(0)}$ and $w_1, \ldots, w_k \in \mathcal{T}^{(0)}(w_0)$. One may assume that $\tau(w_0)$ passes through w_1. Then we must have $\tau(w_1) = \tau(w_0)$ and we choose $\tau(w_i)$ arbitrarily for the other

vertices. Let $m_i := \varphi(\chi_{w_i})$ for $i = 0, \ldots, k$. Then, however we extend τ further and whatever the choice of horizontal edges of level $n+1$ we have $\varphi_{\tau,\mathcal{H}}(w_0) = \varphi(w_0) = m_0$. Now $\mathcal{H}^+_{n+1}$ should consist of $-m_0 + \sum_{i=1}^{k-1} m_i$ edges from w_{k-1} to w_k. Since $\sum_{i=1}^{k} m_i = m_0$ we have $\varphi_{\tau,\mathcal{H}}(\chi_{w_i}) = \varphi(\chi_{w_i})$ for the other i as well.

The construction of the last paragraph is iterated for all vertices at level n, and then for all levels. □

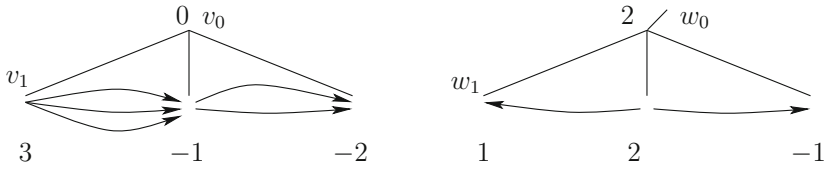

FIGURE 7. The left picture shows the beginning of the tree with its root vertex v_0. The right figure shows a part of the tree corresponding to a branching at vertex w_0 of some higher level. In both cases we assume that the choice function applied to the top vertex corresponds to a path via the left vertical edge. The numbers assigned to vertices stand for the values that φ assigns to the indicator functions corresponding to the vertex. The horizontal edges are chosen according to the rule explained in the proof of Proposition 7.5.

Remark 7.6. It could happen that the construction above provides a Fredholm operator for which some of the $\mathcal{H}_n$ are empty. It still defines a Fredholm module, and adding appropriate weights defines a spectral triple. A Fredholm module can also be modified in order to augment the sets $\mathcal{H}$ without changing the pairing map: given two vertices v and v', simply add an edge from v to v' and an edge from v' to v.

References

[1] J.E. Anderson, I.F. Putnam. "Topological invariants for substitution tilings and their associated C^*-algebras." Ergodic Theory and Dynamical Systems **18** (1998), 509–538.

[2] J. Bellissard. *K-theory of C^*-Algebras in solid state physics.* Statistical mechanics and field theory: mathematical aspects. Springer Berlin Heidelberg, 1986, 99–156.

[3] J. Bellissard, R. Benedetti, J.-M. Gambaudo. "Spaces of Tilings, Finite Telescopic Approximations and Gap-labelling". *Commun. Math. Phys.* **261** (2006), 1–41.

[4] J. Bellissard, A. van Elst, H. Schulz-Baldes. "The noncommutative geometry of the quantum Hall effect." *J. Math. Phys.* **35** (1994), 5373–5451.

[5] J. Bellissard, D. Herrmann, M. Zarrouati. "Hulls of aperiodic solids and gap labeling theorems." In *Directions in mathematical quasicrystals*, volume 13 of *CRM Monogr. Ser.*, pages 207–258. Amer. Math. Soc., Providence, RI, 2000.

[6] J. Bellissard, A. Julien, J. Savinien. "Tiling groupoids and Bratteli diagrams." *Ann. Inst. Henri Poincaré* **11** (2010), 69–99.

[7] J. Bellissard, M. Marcolli, K. Reihani. "Dynamical Systems on Spectral Metric Spaces." *arXiv:1008.4617*.

[8] J. Bellissard, H. Schulz-Baldes. "Scattering theory for lattice operators in dimension $d \geq 3$." Reviews in Mathematical Physics 24.08 (2012).

[9] A. Besbes, D. Lenz. "Delone sets with finite local complexity: linear repetitivity versus positivity of weights." Epreprint available at http://www.math.uni-bielefeld.de/bibos/preprints/08-04-289.pdf

[10] M. Bucci, A. De Luca, A. Glen, L. Zamboni."A new characteristic property of rich words." *Theoret. Comput. Sci.* **410** (2009), 2860–2863.

[11] A. Connes. *Géométrie non commutative*. Paris: InterEditions, 1990.

[12] A. Connes. *Noncommutative geometry*. Academic Press Inc., San Diego, CA, 1994.

[13] E. Christensen, C. Ivan. "Sums of two-dimensional spectral triples". *Math. Scand.* **100** (2007), 35–60.

[14] E. Christensen, C. Ivan, M.L. Lapidus. "Dirac operators and spectral triples for some fractal sets built on curves." *Adv. Math.* **217** (2008), no. 1, 42–78.

[15] J. Cuntz, W. Krieger. "A class of C^*-algebras and topological Markov chains." *Invent. Math.* **56** (1980), 251–268.

[16] D. Damanik, D. Lenz. "Powers in Sturmian sequences". *European Journal of Combinatorics* **24** (2003) 377–390.

[17] F. Durand. "Linearly recurrent subshifts have a finite number of non-periodic subshift factors." *Ergod. Th. & Dynam. Sys.* **20** (2000), 1061–1078.

[18] F. Durand. Corrigendum and addendum to [17]. *Ergo. Th. & Dynam. Sys.* **23** (2003) 663–669.

[19] F. Durand. "Combinatorics on Bratteli diagrams and dynamical systems." *Combinatorics, Automata and Number Theory*, 338–386, Encyclopedia Math. Appl. 135, Cambridge Univ. Press, 2010.

[20] F. Durand, B. Host, C. Skau. "Substitution dynamical systems, Bratteli diagrams and dimension groups." *Ergod. Th. & Dynam. Sys.* **19** (1999), 953–993.

[21] A. Glen, J. Justin, S. Widmer, L.Q. Zamboni. "Palindromic richness." *European J. Combin.* **30** (2009), 510–531.

[22] J. Gracia-Bondia, H. Figueroa, J. Varilly. *Elements of Non-commutative Geometry*. Birkhäuser (2000).

[23] D. Guido, T. Isola. "Dimensions and singular traces for spectral triples, with applications for fractals." *J. Func. Anal.* **203** (2003), 362–400.

[24] D. Guido, T. Isola. "Dimension and spectral triples for fractals in $\mathbb{R}^n$." In Advances in Operator Algebras and Mathematical Physics, *Theta Ser. Adv. Math.* **5**, Theta, Bucarest (2005), 89–108.

[25] B. Grünbaum, G.C. Shephard. *Tilings and patterns*. WH Freeman & Co., 1986.

[26] M. Herman, I.F. Putnam, C. Skau. "Ordered Bratteli diagrams, dimension groups and topological dynamics." *Internat. J. Math.* **3** (1992), 827–864.

[27] A. Julien. "Complexité des pavages apériodiques: calculs et interprétations." PhD Dissertation, Université de Lyon, Université Claude Bernard Lyon 1, 2009.

[28] A. Julien, J. Savinien. "Transverse Laplacians for substitution tilings." *Commun. Math. Phys.* **301** (2010), 285–318.

[29] A. Julien, J. Savinien. "Embeddings of self-similar ultrametric Cantor sets." *Topol. Appl.* **158** (2011), 2148–2157.

[30] J. Kellendonk. "Noncommutative geometry of tilings and gap labelling." *Rev. Math. Phys.* **7** (1995) 1133–1180.

[31] J. Kellendonk. "Topological equivalence of tilings." *J. Math. Phys.* **38** (1997), 1823–1842.

[32] J. Kellendonk. "Gap labelling and the pressure on the boundary." *Commun. Math. Phys.* **258** (2005), 751–768.

[33] J. Kellendonk, D. Lenz, J. Savinien. "A characterization of subshifts with bounded powers." To appear in *Discrete Math.* (2013). Eprint arXiv:1111.1609 (math.CO).

[34] J. Kellendonk, I. Putnam. "Tilings, C^*-algebras and K-theory." In *Directions in mathematical quasicrystals*, volume 13 of *CRM Monogr. Ser.*, pages 177–206. Amer. Math. Soc., Providence, RI, 2000.

[35] J. Kellendonk, S. Richard. "Levinson's theorem for Schrödinger operators with point interaction: a topological approach." *J. Phys. A* **39** (2006).

[36] J. Kellendonk, S. Richard. "The topological meaning of Levinson's theorem, half-bound states included." *J. Phys. A* **41** (2008).

[37] J. Kellendonk, S. Richard. "On the wave operators and Levinson's theorem for potential scattering in $\mathbb{R}^3$." *Asian-Eur. J. Math.* **5** (2012).

[38] J. Kellendonk, T. Richter, H. Schulz-Baldes. "Edge current channels and Chern numbers in the integer quantum Hall effect." *Rev. Math. Phys.* **14** (2002), 87–119.

[39] J. Kellendonk, J. Savinien. "Spectral triples and characterization of aperiodic order." *Proc. London Math. Soc.* **104** (2012) 123–157.

[40] J. Kellendonk, J. Savinien. "Spectral triples from stationary Bratteli diagrams." arXiv:1210.7360.

[41] J. Lagarias, P. Pleasants. "Repetitive Delone sets and quasicrystals." *Ergodic Theory Dynam. Systems* **23** (2003), 831–867.

[42] D. Lenz. "Hierarchical structures in Sturmian dynamical systems. Tilings of the plane." *Theoret. Comput. Sci.* **303** (2003) 463–490.

[43] D. Lenz. "Aperiodic linearly repetitive Delone sets are densely repetitive." *Discrete Comput. Geom.* **31** (2004) 323–326.

[44] G. Michon. "Les Cantors réguliers." *C.-R. Acad. Sci. Paris Sér. I Math.* **300** (1985), 673–675.

[45] H. Moustafa. "PV Cohomology of the Pinwheel Tilings, their Integer Group of Coinvariants and Gap-Labeling." *Commun. Math. Phys.* **298** (2010), 369–405.

[46] E. Prodan, H. Schulz-Baldes, "Non-commutative odd Chern numbers and topological phases of disordered chiral systems." arXiv:1402.5002.

[47] H. Oyono-Oyono, S. Petite. "C^*-algebras of Penrose hyperbolic tilings." *J. Geom. Phys.* **61** (2011), 400–424.

[48] I.C. Palmer. "Riemannian geometry of compact metric spaces." PhD Dissertation, Georgia Institute of Technology 2010.

[49] J. Pearson. "Noncommutative Riemannian geometry and diffusion on ultrametric Cantor sets." PhD Dissertation, Georgia Institute of Technology 2008.

[50] J. Pearson, J. Bellissard. "Noncommutative Riemannian geometry and diffusion on ultrametric Cantor sets." *J. Noncommut. Geom.* **3** (2009), 447–481.

[51] M. Ramirez-Solano. "Non-Commutative Geometrical Aspects And Topological Invariants Of A Conformally Regular Pentagonal Tiling Of The Plane." PhD thesis, Department of Mathematical Sciences, University of Copenhagen, 2013.

[52] M. Ramirez-Solano. "Cohomology of the continuous hull of a combinatorial pentagonal tiling." arXiv preprint arXiv:1305.1062 (2013).

[53] M. Rieffel. "Metrics on state spaces." *Doc. Math.* **4** (1999), 559–600.

[54] M. Rieffel. "Compact Quantum Metric Spaces." *Operator algebras, quantization, and noncommutative geometry* 315–330, Contemp. Math. **365**, Amer. Math. Soc., Providence (2004).

[55] D.J. Rudolph, Rectangular tilings of R^n and free R^n-actions, *Dynamical systems* (College Park, MD, 198–87), Springer, Berlin, 1988, pp. 653–688.

[56] M. Waldschmidt, *private communication.*

[57] M.F. Whittaker. "Spectral triples for hyperbolic dynamical systems." *J. Noncommut. Geom.* **7** (2013), no. 2, 563–582.

Antoine Julien
NTNU
Trondheim, Norway
e-mail: antoine.julien@math.ntnu.no

Johannes Kellendonk
Université Claude Bernard Lyon 1
CNRS, UMR 5208 Institut Camille Jordan
F-69622 Villeurbanne cedex, France
e-mail: kellendonk@math.univ-lyon1.fr

Jean Savinien
Institut Elie Cartan de Lorraine
Université de Lorraine
Metz, France
e-mail: jean.savinien@univ-lorraine.fr

Progress in Mathematics, Vol. 309, 307–370
© 2015 Springer Basel

Spectral Properties of Schrödinger Operators Arising in the Study of Quasicrystals

David Damanik, Mark Embree and Anton Gorodetski

Abstract. We survey results that have been obtained for self-adjoint operators, and especially Schrödinger operators, associated with mathematical models of quasicrystals. After presenting general results that hold in arbitrary dimensions, we focus our attention on the one-dimensional case, and in particular on several key examples. The most prominent of these is the Fibonacci Hamiltonian, for which much is known by now and to which an entire section is devoted here. Other examples that are discussed in detail are given by the more general class of Schrödinger operators with Sturmian potentials. We put some emphasis on the methods that have been introduced quite recently in the study of these operators, many of them coming from hyperbolic dynamics. We conclude with a multitude of numerical calculations that illustrate the validity of the known rigorous results and suggest conjectures for further exploration.

Mathematics Subject Classification (2010). 81Q10, 47B36, 47N50, 35Q41.

Keywords. Schrödinger operators, spectrum, density of states, quantum transport.

1. Introduction

The area of mathematical quasicrystals is fascinating due to its richness and very broad scope. One perspective that leads to a rich theory is the question of how well quantum wave packets can travel in a quasicrystalline medium. This, in turn, leads to a study of the time-dependent Schrödinger equation governed by a potential that reflects the aperiodic order of the environment to which the quantum state is exposed.

D.D. was supported in part by a Simons Fellowship and NSF grants DMS–0800100 and DMS–1067988. M.E. was supported in part by NSF grant DMS–CAREER–0449973. A.G. was supported in part by NSF grants DMS-1301515 and IIS–1018433.

This survey paper describes the current "state of the art" of mathematical results concerning quantum transport in mathematical quasicrystal models. We will describe the (Schrödinger) operators that are typically considered in this context, as well as the spectral and quantum dynamical properties of these operators that are relevant to our basic motivating question. Few results hold without further assumptions. After presenting these general results, we will therefore specialize our class of operators in various ways. On the one hand, additional general results follow from passing to a one-dimensional setting, most notably some specific consequences of Kotani theory. On the other hand, there is much interest in certain central examples, such as the Fibonacci model, and more generally the Sturmian case or potentials generated by primitive substitutions. For these special cases, many more results are known, most of which will be described in detail here. Another purpose of this survey is to highlight new tools that have recently been introduced in the study of these models that have led to very fine quantitative results in the Fibonacci case and beyond.

The structure of the paper is as follows. In Section 2 we describe the operators that are typically considered in the case of general dimension. In Section 3 we present the results known to hold in this very general setting. Next, in Section 4, we pass to the one-dimensional situation, where the operators are slightly redefined to conform with the bulk of the literature. Additional general results are described in this scenario. These range from Kotani theory, via proofs of zero-measure spectrum and hence absence of absolutely continuous spectrum based on Kotani theory, to proofs of absence of point spectrum and hence purely singular continuous spectrum based on Gordon's lemma. We also discuss how these general results can be applied to several classes of examples. Section 5 is devoted to the special case of the Fibonacci Hamiltonian. This is the most prominent one-dimensional quasicrystal model; in addition to the general results mentioned before, there are fine estimates of the local and global dimensions of the spectrum and the density of states measure, as well as the optimal Hölder exponent of the integrated density of states and the upper transport exponent. For this model the question about quantum transport behavior also has very satisfactory answers. All these results, along with comments on their proofs, are presented in this section. Then, in Section 6, we discuss how the approach and the results may be extended from the Fibonacci case to the more general Sturmian case. We present and discuss numerical results in Section 7. Finally, Section 8 lists and discusses a number of open problems that are suggested by the existing results.

Related matters have been surveyed earlier in [11, 27, 43, 46, 166].

2. Schrödinger operators arising in the study of quasicrystals

In this section we describe self-adjoint operators that have been studied in the context of quasicrystals. There is a clear distinction between the case of one space dimension and the case of higher space dimensions. While the geometry of the

quasicrystal model plays an important role in higher dimensions, this is not the case for one-dimensional models. In the former case, this leads to a dependence of the underlying Hilbert space on the realization of the model. In particular, as one typically embeds any such realization in a family of realizations, this leads to some technical issues that need to be addressed mathematically. In the latter case, on the other hand, one usually works in a universal Hilbert space and the aperiodic order features are solely reflected by the potential of the operator. This allows one to invoke the standard theory of ergodic Schrödinger operators with a fixed Hilbert space. In this survey we will present the known results in the settings in which they have been obtained, which we now describe.

Quasicrystals are commonly modeled either by Delone sets in Euclidean space or by tilings of Euclidean space. In fact, any such Delone set or tiling is embedded in a Delone or tiling dynamical system, which is obtained by considering the set of translates and then taking the closure of this set with respect to a suitable topology. This orbit closure is called the hull of the initial Delone set or tiling. The dynamics are then given by the natural action of the Euclidean space on the hull through translations. For this dynamical system, one then identifies ergodic measures, and they are typically unique. While each element of the hull gives rise to a Schrödinger operator, it is the ergodic framework that allows one to prove statements that hold for almost all such operators, as opposed to results for a single such operator. Occasionally, it is then even possible to extend results that hold for almost all elements of the hull to all elements of the hull by approximation.

Delone sets and tilings are in some sense dual and hence equivalent to each other.[1] For definiteness, we will consider a framework based on Delone sets. Let us fix a dimension $d \geq 1$. The Euclidean norm on $\mathbb{R}^d$ will be denoted by $|\cdot|$. We denote by $B(x,r)$ the closed ball in $\mathbb{R}^d$ that is centered at x and has radius r.

Definition 2.1. A set $\Lambda \subset \mathbb{R}^d$ is called a Delone set if there are $r, R > 0$ such that

$$B(x,r) \cap \Lambda = \{x\} \quad \forall x \in \Lambda$$

and

$$B(x,R) \cap \Lambda \neq \emptyset \qquad \forall x \in \mathbb{R}^d.$$

Thus for a Delone set we have a lower bound on the separation between points of the set, and an upper bound for the size of "holes" in the set. One also says that a Delone set is uniformly discrete and relatively dense. Since Delone sets are closed and we want to take orbit closures, let us now define the underlying topology on $\mathcal{F}(\mathbb{R}^d)$, the closed subsets of $\mathbb{R}^d$, following [122].

Definition 2.2. For $k \in \mathbb{Z}_+$ and $F, G \in \mathcal{F}(\mathbb{R}^d)$, we define

$$d_k(F,G) := \inf\left(\{\varepsilon > 0 : F \cap B(0,k) \subset U_\varepsilon(G) \text{ and } G \cap B(0,k) \subset U_\varepsilon(F)\} \cup \{1\}\right),$$

[1] One can go back and forth between these two settings by decorations and the Voronoi construction.

where $U_\varepsilon(\cdot)$ is an open neighborhood. For $k \in \mathbb{Z}_+$, $\varepsilon > 0$, and $F \in \mathcal{F}(\mathbb{R}^d)$, we define

$$U_{\varepsilon,k}(F) := \{G \in \mathcal{F}(\mathbb{R}^d) : d_k(F,G) < \varepsilon\}.$$

The topology on $\mathcal{F}(\mathbb{R}^d)$ with neighborhood basis $\{U_{\varepsilon,k}(F)\}$ will be called the natural topology and denoted by τ_{nat}.

Proposition 2.3.

(a) *Translations are continuous with respect to* τ_{nat}.
(b) $\mathcal{F}(\mathbb{R}^d)$ *endowed with* τ_{nat} *is compact.*
(c) τ_{nat} *is metrizable.*

See [122] for a metric that induces τ_{nat}.

Definition 2.4. Let Ω be a set of Delone sets and denote by T the translation action of $\mathbb{R}^d$, that is, $T_t x = x + t$. The pair (Ω, T) is a *Delone dynamical system* if Ω is invariant under T and closed in the natural topology.

Definition 2.5. A Delone dynamical system is said to be of *finite local complexity* if for every radius $s > 0$, there is a uniform upper bound on the number of different patterns one can observe in ω intersected with a ball of radius s. Here, ω ranges over Ω, the center of the ball ranges over $\mathbb{R}^d$. (A pattern that appears in ω is any finite subset of ω modulo translations.)

Definition 2.6. Suppose (Ω, T) is a Delone dynamical system of finite local complexity. A family $\{A_\omega\}_{\omega\in\Omega}$ of bounded operators $A_\omega : \ell^2(\omega) \to \ell^2(\omega)$ is said to have *finite range* if there is $s > 0$ such that for all $\omega \in \Omega$, $A_\omega(x,y)$ only depends on the pattern of ω in the s-neighborhood of x and y, and $A_\omega(x,y) = 0$ whenever $|x - y| \geq s$.

The class of operators so defined encompasses all discrete operators that are usually considered as quantum Hamiltonians in the context of multi-dimensional quasicrystals. In one dimension, however, it is customary to realign points as a lattice (i.e., $\mathbb{Z}$) and to encode the geometry in the matrix elements of the operator. Even more specifically, one focuses on nearest neighbor interactions and hence obtains a tridiagonal matrix in the standard basis of $\ell^2(\mathbb{Z})$. Much of the mathematical literature focuses on the discrete Schrödinger case, where the terms on the first off-diagonals are all equal to one and the quasicrystalline structure of the environment is reflected in the terms on the diagonal. That is, the family of discrete Schrödinger operators one then considers is of the form $\{H_\omega\}_{\omega\in\Omega}$, where Ω is typically a subshift over a finite alphabet, and for $\omega \in \Omega$, H_ω acts on vectors from $\ell^2(\mathbb{Z})$ as

$$[H_\omega \psi](n) = \psi(n+1) + \psi(n-1) + V_\omega(n)\psi(n).$$

Here the potential V_ω is given by $V_\omega(n) = f(T^n\omega)$, where $T : \Omega \to \Omega$ is the standard shift transformation and f is at least continuous, usually locally constant, and often just depends on a single entry of the sequence ω.

3. General results in arbitrary dimension

3.1. Spectrum and spectral types

One of the fundamental results for the families of operators introduced above is the almost sure constancy of the spectrum and the spectral type with respect to an ergodic measure μ associated with (Ω, T). This follows from the covariance condition

$$A_{\omega+t} = U_t A_\omega U_t^*, \quad \omega \in \Omega,\ t \in \mathbb{R}^d,$$

where $U_t : \ell^2(\omega) \to \ell^2(\omega + t)$ is the unitary operator induced by translation by t, along with the definition of ergodicity applied to traces of spectral projections in the usual way. This establishes the following result; compare [122].

Theorem 3.1. *Suppose (Ω, T) is a Delone dynamical system of finite local complexity and μ is an ergodic measure. Let $\{A_\omega\}_{\omega\in\Omega}$ be a family of bounded self-adjoint operators of finite range. Then there exist $\Sigma, \Sigma_{\rm pp}, \Sigma_{\rm sc}, \Sigma_{\rm ac}$ and a subset $\Omega_0 \subseteq \Omega$ of full μ-measure such that for every $\omega \in \Omega_0$, we have $\sigma(A_\omega) = \Sigma$, $\sigma_{\rm pp}(A_\omega) = \Sigma_{\rm pp}$, $\sigma_{\rm sc}(A_\omega) = \Sigma_{\rm sc}$, and $\sigma_{\rm ac}(A_\omega) = \Sigma_{\rm ac}$.*

3.2. Existence of the integrated density of states

Suppose $\{A_\omega\}_{\omega\in\Omega}$ is a family of bounded self-adjoint operators of finite range. For $Q \subset \mathbb{R}^d$ bounded, the restriction $A_\omega|_Q$ defined on $\ell^2(Q \cap \omega)$ has finite rank. Therefore,

$$n(A_\omega, Q)(E) := \#\{\text{eigenvalues of } A_\omega|Q \text{ that are } \le E\}$$

is finite and $E \mapsto \frac{1}{|Q|} n(A_\omega, Q)(E)$ is the distribution function of the measure $\rho_Q^{A_\omega}$ defined by

$$\int \varphi \, d\rho_Q^{A_\omega} = \frac{1}{|Q|} \mathrm{Tr}(\varphi(A_\omega|Q)), \quad \varphi \in C_b(\mathbb{R}).$$

For $s > 0$ and $Q \subseteq \mathbb{R}^d$, denote by $\partial_s Q$ the set of points in $\mathbb{R}^d$ whose distance from the boundary of Q is bounded by s. A sequence $\{Q_k\}$ of bounded subsets of $\mathbb{R}^d$ is called a van Hove sequence if $\frac{\mathrm{vol}(\partial_s Q_k)}{\mathrm{vol}(Q_k)} \to 0$ as $k \to \infty$ for every $s > 0$.

The following result was shown in [123] (compare also the earlier papers [93, 94]).

Theorem 3.2. *Suppose (Ω, T) is a strictly ergodic[2] Delone dynamical system of finite local complexity. Let $\{A_\omega\}_{\omega\in\Omega}$ be a family of bounded self-adjoint operators of finite range and let $\{Q_k\}$ be a van Hove sequence. Then, as $k \to \infty$, the distributions of $E \mapsto \rho_{Q_k}^{A_\omega}((-\infty, E])$ converge to the distribution $E \mapsto \rho^A((-\infty, E])$ with respect to $\|\cdot\|_\infty$, and the convergence is uniform in $\omega \in \Omega$.*

In fact, the limiting distribution can be given in closed form; see [123]. It is called the integrated density of states, and the associated measure is called the density of states measure. The remarkable feature of this result is the strength

[2]Strict ergodicity means that all orbits are dense and that there is a unique invariant Borel probability measure.

of the convergence, in that the distribution functions converge uniformly in the $\|\cdot\|_\infty$ topology. This is of particular interest in cases when the limiting distribution function has jumps. The next subsection shows that the latter phenomenon may actually happen.

3.3. Locally supported eigenfunctions and discontinuities of the IDS

In dimensions strictly greater than one, the local structure of a Delone set may be chosen such that suitable finite range operators have finitely supported eigenfunctions at a suitable energy. If these local configurations occur sufficiently regularly, it follows that the energy in question will be a point of discontinuity of the integrated density of states. This observation may now be supplemented in two ways. On the one hand, a given Delone dynamical system may be transformed into one that is equivalent to the original one in the sense of mutual local derivability, which does have the required local configurations. On the other hand, any discontinuity of the integrated density of states must arise in this way, that is, through the regular occurrence of finitely supported eigenfunctions at the energy in question. These issues were discussed in the paper [111] (see that paper for the definition of mutual local derivability). Let us state the results from that paper precisely.

Theorem 3.3. *Suppose (Ω, T) is a strictly ergodic Delone dynamical system of finite local complexity. Let $\{A_\omega\}_{\omega\in\Omega}$ be a family of bounded self-adjoint operators of finite range. Then there exist (Ω', T) and $\{A'_\omega\}_{\omega\in\Omega'}$ such that (Ω, T) and (Ω', T) are mutually locally derivable and A'_ω has locally supported eigenfunctions with the same eigenvalue for every $\omega \in \Omega'$. Moreover, A'_ω can be chosen to be the nearest neighbor Laplacian of a suitable graph.*

Theorem 3.4. *Suppose (Ω, T) is a strictly ergodic Delone dynamical system of finite local complexity. Let $\{A_\omega\}_{\omega\in\Omega}$ be a family of bounded self-adjoint operators of finite range. Then $E \in \mathbb{R}$ is a point of discontinuity of ρ^A if and only if there exists a locally supported eigenfunction of A_ω for E for one (equivalently, all) $\omega \in \Omega$.*

4. General results in one dimension

Starting with this section, we will focus on the case of one space dimension. Far more rigorous results are known for this special case than for the general case. In particular, much is known about the structure of the spectrum as a set, as well as the type of the spectral measures.

As mentioned above, in the one-dimensional setting one typically passes to a somewhat different choice of the model. Thus, for definiteness, we will restrict our attention in much of the remainder of this paper to the following scenario. We consider Schrödinger operators in $\ell^2(\mathbb{Z})$,

$$[H_\omega\psi](n) = \psi(n+1) + \psi(n-1) + V_\omega(n)\psi(n), \tag{1}$$

where the potentials are of the form $V_\omega(n) = f(T^n\omega)$, with ω in a compact metric space Ω, a homeomorphism $T : \Omega \to \Omega$, and $f \in C(\Omega, \mathbb{R})$. We also fix an ergodic measure μ.

Notice that the Hilbert space in which H_ω acts is now ω-independent, and the aperiodic order features of the medium that is being modeled are completely subsumed in the potential V_ω of the operator H_ω. In this we follow the standard convention, for this class of operators has been commonly studied. One could consider operators that are formally more akin to the operators considered above in general dimensions. However, this would not lead to any significant mathematical difference. Loosely speaking, the aperiodically-ordered Delone set in $\mathbb{R}$ is just being reconfigured as $\mathbb{Z}$, and the local properties of the operator that depend on the pattern near a point in the general setting affect the value of the potential at the point in question accordingly in our present setting.

In fact, this scenario is more general than considered in the previous section. The better analog would be the case where $(T\omega)_n = \omega_{n+1}$ is the shift on $\mathcal{A}^{\mathbb{Z}}$ for some finite set $\mathcal{A}$, $\Omega \subseteq \mathcal{A}^{\mathbb{Z}}$ is T-invariant and closed, and $f : \Omega \to \mathbb{R}$ is locally constant, that is, it only depends on the values of ω_n for some finite set of n values. However, some of the results below hold in the more general setting, and we will impose further restrictions when they are needed.

Consequently, Theorems 3.1 and 3.2 now take the following form.

Theorem 4.1. *There are sets Σ, Σ_{ac}, Σ_{sc}, and Σ_{pp} such that for μ-almost every $\omega \in \Omega$, we have $\sigma(H_\omega) = \Sigma$, $\sigma_{\mathrm{ac}}(H_\omega) = \Sigma_{\mathrm{ac}}$, $\sigma_{\mathrm{sc}}(H_\omega) = \Sigma_{\mathrm{sc}}$, and $\sigma_{\mathrm{pp}}(H_\omega) = \Sigma_{\mathrm{pp}}$.*

Theorem 4.2. *The measures $\int \varphi \, dN_k^\omega = \frac{1}{k}\mathrm{Tr}(\varphi(H_\omega|_{[1,k]}))$ converge weakly to the measure $\int \varphi \, dN = \int \langle \delta_0, \varphi(H_\omega)\delta_0 \rangle \, d\mu(\omega)$ as $k \to \infty$.*

The second result uses a weaker notion of convergence than in Theorem 3.2, the price we have to pay for casting this problem in the more general setting. However, this is fine after all, due to the following result.

Theorem 4.3. *The measure dN is continuous.*

Moreover, we have:

Theorem 4.4. *The topological support of the measure dN is equal to Σ.*

Theorems 4.1–4.4 are standard results from the theory of ergodic Schrödinger operators on $\ell^2(\mathbb{Z})$; compare [32].

4.1. Spectrum and the absence of uniform hyperbolicity

Let us consider solutions to the difference equation

$$u(n+1) + u(n-1) + V_\omega(n)u(n) = Eu(n), \quad n \in \mathbb{Z}, \tag{2}$$

for some energy $E \in \mathbb{R}$. Clearly, u solves (2) if and only if

$$\begin{pmatrix} u(n+1) \\ u(n) \end{pmatrix} = \begin{pmatrix} E - V_\omega(n) & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} u(n) \\ u(n-1) \end{pmatrix}, \quad n \in \mathbb{Z}. \tag{3}$$

Since $V_\omega(n) = f(T^n\omega)$, one naturally defines

$$A_E(\omega) = \begin{pmatrix} E - f(\omega) & -1 \\ 1 & 0 \end{pmatrix},$$

so that (3) implies

$$\begin{pmatrix} u(n+1) \\ u(n) \end{pmatrix} = A_E(T^n\omega) \times \cdots \times A_E(T\omega) \begin{pmatrix} u(1) \\ u(0) \end{pmatrix}$$

for $n \geq 1$ and solutions u to (2). We set $M_{E,\omega}(n) = A_E(T^n\omega) \times \cdots \times A_E(T\omega)$.

Definition 4.5. We let

$$\mathcal{UH} = \{E \in \mathbb{R} : \exists\, c > 1 \text{ such that } \forall \omega \in \Omega,\ n \geq 1 \text{ we have } \|M_{E,\omega}(n)\| \geq c^n\}.$$

There are a number of equivalent ways to describe the uniform hyperbolicity of $M_{E,\omega}(n)$, such as an invariant splitting into stable and unstable directions and the absence of a Sacker-Sell solution; compare, for example, [49, 178, 179].

Johnson showed in [105] that the set $\mathcal{UH}$ is equal to the resolvent set of H_ω for any ω that has a dense T-orbit. As a particular consequence, we may state the following:

Theorem 4.6. *Suppose T is minimal. Then for every $\omega \in \Omega$, we have $\sigma(H_\omega) = \mathbb{R} \setminus \mathcal{UH}$. In particular, for any ergodic measure μ, we have $\Sigma = \mathbb{R} \setminus \mathcal{UH}$.*

4.2. Kotani theory

In the previous subsection, we saw that the partition $\mathbb{R} = \mathcal{UH} \sqcup \mathbb{R} \setminus \mathcal{UH}$ yields the partition of the energy axis into resolvent set and spectrum. Let us partition the energy axis even further by introducing the Lyapunov exponent

$$L_\mu(E) = \lim_{n\to\infty} \frac{1}{n} \int \log \|M_{E,\omega}(n)\|\, d\mu(\omega).$$

The existence of the limit follows quickly by subadditivity. Moreover, due to the ergodicity of μ and Kingman's Subadditive Ergodic Theorem, we have

$$L_\mu(E) = \lim_{n\to\infty} \frac{1}{n} \log \|M_{E,\omega}(n)\| \quad \text{for } \mu\text{-a.e. } \omega \in \Omega.$$

Obviously, we have $L_\mu(E) > 0$ for every $E \in \mathcal{UH}$. We let

$$\mathcal{Z}_\mu = \{E : L_\mu(E) = 0\}$$

and

$$\mathcal{NUH}_\mu = \{E \in \mathbb{R} : L_\mu(E) > 0\} \setminus \mathcal{UH},$$

so that our final partition is $\mathbb{R} = \mathcal{UH} \sqcup \mathcal{NUH}_\mu \sqcup \mathcal{Z}_\mu$. Notice that in these definitions, the dependence of the partition of $\mathbb{R} \setminus \mathcal{UH}$ into $\mathcal{NUH}_\mu \sqcup \mathcal{Z}_\mu$ on the choice of the ergodic measure μ is emphasized through the subscript for the μ-dependent sets. It is customary to drop this explicit subscript from L, $\mathcal{Z}$, and $\mathcal{NUH}$, and we will henceforth do so as well.

Recall that the essential closure of a set $S \subseteq \mathbb{R}$ is given by

$$\overline{S}^{\text{ess}} = \{E \in \mathbb{R} : \forall\, \varepsilon > 0 \text{ we have } \mathrm{Leb}(S \cap (E - \varepsilon, E + \varepsilon)) > 0\}.$$

Then, the following result is the celebrated Ishii–Pastur–Kotani theorem; see, for example, [47, 102, 115, 144, 161].

Theorem 4.7. *We have* $\Sigma_{\mathrm{ac}} = \overline{\mathcal{Z}}^{\mathrm{ess}}$.

Since the essential closure of a set S is empty if and only if S has zero Lebesgue measure, this result yields a characterization of the almost sure purely singular spectrum. In fact, this is the typical situation for our models of interest [116].

Theorem 4.8. *Suppose the range of* $f : \Omega \to \mathbb{R}$ *is finite and* $\mathrm{Leb}(\mathcal{Z}) > 0$. *Denote the push-forward of* μ *under* $\Omega \ni \omega \mapsto V_\omega \in (\mathrm{Ran}\, f)^{\mathbb{Z}}$ *by* μ^*. *Then,* $\mathrm{supp}\, \mu^*$ *is finite.*

Since $\mathrm{supp}\, \mu$ is T-invariant, this means that every element of $\mathrm{supp}\, \mu^*$ is periodic. In other words, if the potentials are ergodic, aperiodic and take finitely many values, then $\mathcal{Z}$ has zero Lebesgue measure and the almost sure absolutely continuous spectrum is empty.

4.3. Zero-measure spectrum

The realization that $\mathrm{Leb}(\mathcal{Z}) = 0$ for potentials that are ergodic, aperiodic and take finitely many values is at the heart of proofs of zero-measure spectrum in these cases. Whenever Theorem 4.8 applies, it suffices to show that the spectrum is contained in (and hence coincides with) $\mathcal{Z}$ in order to establish zero-measure spectrum. There are two approaches that establish this identity. One shows directly that the Lyapunov exponent must vanish for every energy in the spectrum by providing subexponential upper bounds for transfer matrix norms. The other is based on Theorem 4.6, which says that $\Sigma = \mathcal{Z} \cup \mathcal{NUH}$. Hence, it suffices to prove that if E is such that the Lyapunov exponent is positive, the convergence of $\frac{1}{n} \log \|M_{E,\omega}(n)\|$ to $L(E)$ must be uniform, as this will imply that $E \in \mathcal{UH}$, and as a consequence, $\Sigma = \mathcal{Z}$. The second approach can be made to work in greater generality, whereas the first approach often gives additional information that has other interesting consequences.

We will discuss how the first approach is implemented when we discuss Sturmian models and, specifically, the Fibonacci Hamiltonian in later sections. Here we therefore discuss in some detail how the second approach works. We first state the main result, then explain how it is a natural consequence of the line of reasoning employed. We emphasize that this is a result that holds in the symbolic setting. That is, for a finite set $\mathcal{A}$, called the *alphabet*, we consider the shift transformation $T : \mathcal{A}^{\mathbb{Z}} \to \mathcal{A}^{\mathbb{Z}}$ given by $(T\omega)(n) = \omega(n+1)$ and closed, T-invariant subsets Ω of $\mathcal{A}^{\mathbb{Z}}$, which are called *subshifts* (*over* $\mathcal{A}$). We denote by $\mathcal{W}_\Omega$ the set of all finite words over the alphabet $\mathcal{A}$ that occur in elements of Ω, and we write $\mathcal{W}_\Omega(n)$ for those elements of $\mathcal{W}_\Omega$ that have length n. A function on Ω is called *locally constant* if $f(\omega)$ only depends on finitely many entries. More formally, there exists $k \in \mathbb{Z}_+$ such that $f(\omega) = f(\omega')$ for all $\omega, \omega' \in \Omega$ with $\omega_{-k} \cdots \omega_k = \omega'_{-k} \cdots \omega'_k$. Note that f is locally constant if and only if it is continuous and takes only finitely many values.

As usual, a subshift Ω is called *minimal* if the topological dynamical system (Ω, T) is minimal (i.e., the T-orbit of every $\omega \in \Omega$ is dense in Ω), it is called *uniquely ergodic* if (Ω, T) has a unique invariant Borel probability measure, and it is called *strictly ergodic* if it is both minimal and uniquely ergodic. In the uniquely ergodic case, the unique invariant Borel probability measure must necessarily be ergodic.

Definition 4.9. Let Ω be a strictly ergodic subshift with unique T-invariant measure μ. It satisfies the Boshernitzan condition (B) if

$$\limsup_{n\to\infty} \left(\min_{w\in\mathcal{W}_\Omega(n)} n \cdot \mu\left([w]\right) \right) > 0.$$

Here is the result from [62] showing that (B) is a sufficient condition for zero-measure spectrum:

Theorem 4.10. *Suppose Ω is a strictly ergodic subshift that satisfies condition* (B) *and $f : \Omega \to \mathbb{R}$ is locally constant. Then the convergence of $\frac{1}{n} \log \|M_{E,\omega}(n)\|$ to $L(E)$ is uniform for every $E \in \mathbb{R}$. In particular, $\mathcal{NUH}$ is empty and $\Sigma = \mathcal{Z}$. Thus, if Ω and f are such that the V_ω are aperiodic, then* $\mathrm{Leb}(\Sigma) = 0$. *In the latter case,* $\Sigma_{\mathrm{ac}} = \emptyset$.

The proof of Theorem 4.10 is more easily understood if one imposes a stronger assumption. Indeed, let us assume for the time being that $f(\omega)$ depends only on ω_0 and $\min_{w\in\mathcal{W}_\Omega(n)} n \cdot \mu\left([w]\right)$ is uniformly large for all n, not merely for a subsequence. That is, suppose there is $\delta > 0$ such that $|w| \cdot \mu\left([w]\right) \geq \delta$ for every $w \in \mathcal{W}_\Omega$.

By our stronger assumption on f, we may view $\log \|M_{E,\omega}(n)\|$ as a quantity associated with the word $w = \omega_1 \cdots \omega_n \in \mathcal{W}_\Omega(n)$ that we will denote by $F(w)$. If we do so, then the goal is to prove that $|w|^{-1} F(w)$ converges uniformly as $|w| \to \infty$ and each w belongs to $\mathcal{W}_\Omega$. It is a known consequence of unique ergodicity of (Ω, T) that the convergence in

$$F^+ := \limsup_{w\in\mathcal{W}_\Omega,\, |w|\to\infty} |w|^{-1} F(w)$$

is uniform. By the uniform upper bound and the frequent occurrence of any word due to the strengthened assumption, one can derive a similar uniform result for the $\liminf$, and hence establish uniform convergence. These realizations are contained in the paper [121] by Lenz.

If one only has condition (B), then a similar way of reasoning can be carried out for the sequence of length scales from (B). In this way, one can establish uniform convergence along this sequence. To interpolate, one can employ the Avalanche Principle from [86].

The paper [63] contains numerous applications of Theorem 4.10 to specific classes of subshifts, some of which will be described in Subsection 4.4 below.

When [62, 63] appeared, all known zero-measure results for Schrödinger operators defined by strictly ergodic subshifts (see, e.g., [10, 12, 14, 26, 120, 121, 129, 165]) were covered by this approach, that is, they all held for subshifts that satisfy

condition (B). Recently, Liu and Qu constructed examples of strictly ergodic subshifts that do not satisfy (B) but for which the associated Schrödinger operators do have zero-measure spectrum (and in fact the convergence of $\frac{1}{n}\log\|M_{E,\omega}(n)\|$ to $L(E)$ is uniform for every $E \in \mathbb{R}$) [127].

Naturally, once one knows that the spectrum has zero Lebesgue measure, one would like to determine its fractal (e.g., Hausdorff, lower and upper box counting) dimensions, as well as similar quantities such as thickness and denseness. These more delicate questions have been studied for a rather small number of examples, which will be discussed in subsequent sections. Zero-measure spectrum, on the other hand, is known in much greater generality, and this is the topic of the next subsection.

4.4. Examples

In this subsection, we present several classes of popular examples of potentials that are ergodic, aperiodic, and take finitely many values (so that Kotani's central result applies) and discuss the validity of condition (B) (so that the associated Schrödinger operators have zero-measure spectrum). For more details, we refer the reader to [63].

4.4.1. Linearly recurrent subshifts and subshifts generated by primitive substitutions. A subshift Ω over $\mathcal{A}$ is called *linearly recurrent* (or *linearly repetitive*) if there exists a constant K such that if $v, w \in \mathcal{W}(\Omega)$ with $|w| \geq K|v|$, then v is a subword of w. Clearly, every linearly recurrent subshift Ω satisfies (B). A popular way to generate linearly recurrent subshifts is via primitive substitutions. A substitution $S : \mathcal{A} \to \mathcal{A}^*$ is called *primitive* if there exists $k \in \mathbb{N}$ such that for every $a, b \in \mathcal{A}$, $S^k(a)$ contains b. Such a substitution generates a subshift Ω as follows. It is easy to see that there are $m \in \mathbb{N}$ and $a \in \mathcal{A}$ such that $S^m(a)$ begins with a. If we iterate S^m on the symbol a, we obtain a one-sided infinite limit, u, called a *substitution sequence*. Ω then consists of all two-sided sequences for which all subwords are also subwords of u. One can verify that this construction is in fact independent of the choice of u, and hence Ω is uniquely determined by S. Prominent examples[3] are given by

$a \mapsto ab,\ b \mapsto a$	Fibonacci
$a \mapsto ab,\ b \mapsto ba$	Thue–Morse
$a \mapsto ab,\ b \mapsto aa$	Period Doubling
$a \mapsto ab,\ b \mapsto ac,\ c \mapsto db,\ d \mapsto dc$	Rudin–Shapiro

The following was shown in [78] (and independently in [72]):

Proposition 4.11. *If the subshift Ω is generated by a primitive substitution, then it is linearly recurrent and hence satisfies condition* (B).

[3]These examples appear explicitly in many papers in the physics literature on Schrödinger operators generated by primitive substitution; compare [2, 6, 15, 27, 84, 96, 114, 131].

It may happen that a non-primitive substitution generates a linearly recurrent subshift. An example is given by $a \mapsto aaba$, $b \mapsto b$. In fact, the class of linearly recurrent subshifts generated by substitutions was characterized in [64].[4] In particular, it turns out that a subshift generated by a substitution is linearly recurrent if and only if it is minimal.

4.4.2. Sturmian and quasi-Sturmian subshifts. Consider a minimal subshift Ω over $\mathcal{A}$. The (factor) complexity function $p : \mathbb{Z}_+ \to \mathbb{Z}_+$ is defined by $p(n) = \#\mathcal{W}_\Omega(n)$. Hedlund and Morse showed in [98] that Ω is aperiodic if and only if $p(n) \geq n+1$ for every $n \in \mathbb{Z}_+$. Aperiodic minimal subshifts of minimal complexity, $p(n) = n+1$ for every $n \in \mathbb{N}$, exist and they are called Sturmian. If the complexity function satisfies $p(n) = n + k$ for $n \geq n_0$, $k, n_0 \in \mathbb{N}$, the subshift is called quasi-Sturmian. It is known that quasi-Sturmian subshifts are exactly those subshifts that are a morphic image of a Sturmian subshift; compare [34, 36, 148].

There are a large number of equivalent characterizations of Sturmian subshifts; compare [18]. We are mainly interested in their geometric description in terms of an irrational rotation. Let $\alpha \in (0,1)$ be irrational and consider the rotation by α on the circle, $R_\alpha : [0,1) \to [0,1)$, $R_\alpha\theta = \{\theta + \alpha\}$, where $\{x\}$ denotes the fractional part of x, $\{x\} = x \mod 1$. The coding of the rotation R_α according to a partition of the circle into two half-open intervals of length α and $1-\alpha$, respectively, is given by the sequences $v_n(\alpha,\theta) = \chi_{[0,\alpha)}(R_\alpha^n\theta)$. We obtain a subshift

$$\Omega_\alpha = \overline{\{v(\alpha,\theta) : \theta \in [0,1)\}} = \{v(\alpha,\theta) : \theta \in [0,1)\} \cup \{\tilde{v}^{(k)}(\alpha) : k \in \mathbb{Z}\} \subset \{0,1\}^{\mathbb{Z}},$$

which can be shown to be Sturmian. Here, $\tilde{v}_n^{(k)}(\alpha) = \chi_{(0,\alpha]}(R_\alpha^{n+k}0)$. Conversely, every Sturmian subshift is essentially of this form, that is, if Ω is minimal and has complexity function $p(n) = n+1$, then, up to a one-to-one morphism, $\Omega = \Omega_\alpha$ for some irrational $\alpha \in (0,1)$.

Using this description and some classical results in diophantine approximation, the following result was shown in [63].

Theorem 4.12. *Every Sturmian subshift obeys the Boshernitzan condition* (B).

Moreover, establishing stability of (B) under morphic images, one obtains the following consequence.

Corollary 4.13. *Every quasi-Sturmian subshift obeys* (B).

4.4.3. Circle maps. Let $\alpha \in (0,1)$ be irrational and $\beta \in (0,1)$ arbitrary. The coding of the rotation R_α according to a partition into two half-open intervals of length β and $1-\beta$, respectively, is given by the sequences $v_n(\alpha,\beta,\theta) = \chi_{[0,\beta)}(R_\alpha^n\theta)$. We obtain a subshift

$$\Omega_{\alpha,\beta} = \overline{\{v(\alpha,\beta,\theta) : \theta \in [0,1)\}} \subset \{0,1\}^{\mathbb{Z}}.$$

[4] See also [75, 76, 126] for results for Schrödinger operators arising from a specific class of non-primitive substitutions.

Subshifts generated this way are usually called circle map subshifts or subshifts generated by the coding of a rotation.

The paper [63] established the following results for circle map subshifts in connection with property (B):

Theorem 4.14. *Let $\alpha \in (0,1)$ be irrational. Then the subshift $\Omega_{\alpha,\beta}$ satisfies* (B) *for Lebesgue almost every $\beta \in (0,1)$.*

Theorem 4.15. *Let $\alpha \in (0,1)$ be irrational with bounded continued fraction coefficients, that is, $a_n \leq C$. Then $\Omega_{\alpha,\beta}$ satisfies* (B) *for every $\beta \in (0,1)$.*

Theorem 4.16. *Let $\alpha \in (0,1)$ be irrational with unbounded continued fraction coefficients. Then there exists $\beta \in (0,1)$ such that $\Omega_{\alpha,\beta}$ does not satisfy* (B).

4.4.4. Interval exchange transformations. Subshifts generated by interval exchange transformations (IETs) are natural generalizations of Sturmian subshifts. IETs are defined as follows. Given a probability vector $\lambda = (\lambda_1, \ldots, \lambda_m)$ with $\lambda_i > 0$ for $1 \leq i \leq m$, we let $\mu_0 = 0$, $\mu_i = \sum_{j=1}^{i} \lambda_j$, and $I_i = [\mu_{i-1}, \mu_i)$. Let τ be a permutation of $\mathcal{A}_m = \{1, \ldots, m\}$, that is, $\tau \in S_m$, the symmetric group. Then $\lambda^\tau = (\lambda_{\tau^{-1}(1)}, \ldots, \lambda_{\tau^{-1}(m)})$ is also a probability vector, and we can form the corresponding μ_i^τ and I_i^τ. Denote the unit interval $[0,1)$ by I. The (λ, τ) interval exchange transformation is then defined by

$$T : I \to I, \;\; T(x) = x - \mu_{i-1} + \mu^\tau_{\tau(i)-1} \text{ for } x \in I_i, \; 1 \leq i \leq m.$$

It exchanges the intervals I_i according to the permutation τ.

The transformation T is invertible, and its inverse is given by the $(\lambda^\tau, \tau^{-1})$ interval exchange transformation.

The symbolic coding of $x \in I$ is $\omega_n(x) = i$ if $T^n(x) \in I_i$. This induces a subshift over the alphabet $\mathcal{A}_m$: $\Omega_{\lambda,\tau} = \overline{\{\omega(x) : x \in I\}}$.

Sturmian subshifts correspond to the case of two intervals, as a first return map construction shows.

Keane [108] proved that if the orbits of the discontinuities μ_i of T are all infinite and pairwise distinct, then T is minimal. In this case, the coding is one-to-one and the subshift is minimal and aperiodic. This holds in particular if τ is irreducible and λ is irrational. Here, τ is called irreducible if $\tau(\{1, \ldots, k\}) \neq \{1, \ldots, k\}$ for every $k < m$ and λ is called irrational if the λ_i are rationally independent.

Regarding property (B), Boshernitzan has proved two results. First, in [21] the following is shown:

Theorem 4.17. *For every irreducible $\tau \in S_m$ and for Lebesgue almost every λ, the subshift $\Omega_{\lambda,\tau}$ satisfies* (B).

In fact, Boshernitzan shows that for every irreducible $\tau \in S_m$ and for Lebesgue almost every λ, the subshift $\Omega_{\lambda,\tau}$ satisfies a stronger condition where the sequence of n values for which $\eta(n)$ is large cannot be too sparse. This condition is easily seen to imply (B), and hence the theorem above.

In a different paper, [22], Boshernitzan singles out an explicit class of subshifts arising from interval exchange transformations that satisfy (B). The transformation T is said to be of (rational) rank k if the μ_i span a k-dimensional space over $\mathbb{Q}$ (the field of rational numbers).

Theorem 4.18. *If T has rank 2, the subshift $\Omega_{\lambda,\tau}$ satisfies* (B).

4.5. Singular continuous spectrum

As seen in the previous section, the spectrum has zero Lebesgue measure whenever condition (B) holds. This condition is satisfied by a wide class of models, in particular by all typical quasicrystal models. As pointed out in Theorem 4.10, a consequence of zero-measure spectrum is the absence of absolutely continuous spectrum. That is, if $\sigma(H_\omega)$ has zero Lebesgue measure, then $\sigma_{\mathrm{ac}}(H_\omega) = \emptyset$, since all spectral measures are supported by the spectrum, and any measure supported by a set of zero Lebesgue measure must be purely singular by definition.

To complement this, one can often show the absence of point spectrum. That is, there are a variety of tools that allow one to show that H_ω has no eigenvalues, and hence $\sigma_{\mathrm{pp}}(H_\omega) = \emptyset$ as well. Putting the two results together, one obtains that H_ω has purely singular continuous spectrum.

The primary tool that allows one to exclude eigenvalues is based on the Gordon lemma, which assumes that the potential has infinitely many suitably aligned local periodicities. Overall, this nicely implements the philosophy that aperiodic order is intermediate between periodic and random. The aperiodicity implies the absence of absolutely continuous spectrum via Kotani's theorem (and hence one does not have the spectral type that appears for a periodic medium), while the order feature implies the absence of point spectrum via a fingerprint of local periodicity (and hence one does not have the spectral type that appears for a random medium).[5]

A potential $V : \mathbb{Z} \to \mathbb{R}$ is called a *Gordon potential* if there are $q_k \to \infty$ such that for every k, we have $V(n) = V(n+q_k) = V(n-q_k)$ for $1 \le n \le q_k$. That is, V looks like a periodic potential around the origin, as one sees at least three suitably aligned periodic unit cells there, and the period may be chosen arbitrarily large. The following Gordon lemma is based in spirit on [88]. In this particular form it was shown in [73].

Lemma 4.19. *Suppose V is a Gordon potential. Then, for every E, the difference equation*

$$u(n+1) + u(n-1) + V(n)u(n) = Eu(n)$$

has no non-trivial square-summable solutions. In particular, the associated Schrödinger operator H in $\ell^2(\mathbb{Z})$, given by

$$[H\psi](n) = \psi(n+1) + \psi(n-1) + V(n)\psi(n),$$

has empty point spectrum.

[5] In a random model, the values of the potential at the various sites are given by independent identically distributed random variables. This model is usually called the Anderson model.

By ergodicity, T-invariance, and the Gordon lemma, if one can show that

$$\mu\left(\{\omega \in \Omega : V_\omega \text{ is a Gordon potential}\}\right) > 0,$$

then

$$\mu\left(\{\omega \in \Omega : H_\omega \text{ has empty point spectrum}\}\right) = 1.$$

On the other hand, for any aperiodic minimal subshift, at least one of its elements fails to have the required Gordon three-block symmetries [42]. Thus, one cannot use this appraoch to show uniform absence of eigenvalues for all $\omega \in \Omega$. Nevertheless, results to this effect are known, established with the following variant of the Gordon lemma.

Lemma 4.20. *Suppose $V : \mathbb{Z} \to \mathbb{R}$ is such that there are $q_k \to \infty$ such that $V(n) = V(n + q_k)$ for $1 \le n \le q_k$. Suppose further that E is such that*

$$\sup_k \left| \mathrm{Tr} \left(\begin{pmatrix} E - V(q_k) & -1 \\ 1 & 0 \end{pmatrix} \times \cdots \times \begin{pmatrix} E - V(1) & -1 \\ 1 & 0 \end{pmatrix} \right) \right| < \infty. \tag{4}$$

Then, the difference equation

$$u(n+1) + u(n-1) + V(n)u(n) = Eu(n)$$

has no non-trivial solutions that are square-summable on $\mathbb{Z}_+$ and hence E is not an eigenvalue of the associated Schrödinger operator H in $\ell^2(\mathbb{Z})$. In particular, if the assumption (4) *holds for every E in the spectrum of H, then H has empty point spectrum.*

In many quasicrystal models, the existence of hierarchical structures gives rise to a so-called trace map, which in turn can often be used to ensure that (4) holds for all energies in the spectrum. Thus, the analysis then reduces to finding suitable squares of arbitrary length starting at the origin.

Thus, in the symbolic setting at hand, the observations above give rise to problems that concern the subword structure of the potentials, and hence fall in the general area of combinatorics on words.

Let us describe the results that have been obtained in this way for the examples discussed above. In all these results, the choice of the sampling function f is more restricted than above. Namely, one usually assumes that $f(\omega) = g(\omega_0)$ with an injective map $g : \mathcal{A} \to \mathbb{R}$.

4.5.1. Subshifts generated by primitive substitutions. Suppose S is a primitive substitution over the alphabet $\mathcal{A}$ and let $\Omega \subseteq \mathcal{A}^{\mathbb{Z}}$ be the subshift associated with it. Recall that it is strictly ergodic and denote the unique invariant probability measure by μ. The *index* of Ω is given by the largest fractional power occurring in an (and hence any) element of Ω. Formally, the index is defined as follows. Given $w \in \mathcal{W}_\Omega$ and any prefix v of w, we denote the word $w^k v$ with $k \in \mathbb{Z}_+$ by w^r, where $r = k + \frac{|v|}{|w|}$. Then, $\mathrm{ind}_\Omega(w) = \sup\{r \in \mathbb{Q} \cap [1, \infty) : w^r \in \mathcal{W}_\Omega\}$ and $\mathrm{ind}(\Omega) = \sup\{\mathrm{ind}_\Omega(w) : w \in \mathcal{W}_\Omega\}$.

The following result was shown in [42].[6]

Theorem 4.21. *If S is a primitive substitution and the associated subshift Ω satisfies* $\operatorname{ind}(\Omega) > 3$*, then* $\mu(\{\omega \in \Omega : V_\omega \text{ is a Gordon potential}\}) > 0$*. Consequently,* $\mu(\{\omega \in \Omega : H_\omega \text{ has empty point spectrum}\}) = 1$.

The underlying idea is simple. Since the subshift is invariant under S, any word appearing with index strictly greater than 3 generates by iteration of S a sequence of words whose lengths go to infinity and whose index is bounded away from 3. This allows one to bound from below the frequency with which third powers occur and hence yields measure estimates on the Gordon three-block conditions that are good enough to show that the lim sup of these sets must have positive measure. Since the elements of the lim sup of these sets give rise to Gordon potentials, the result follows.

Applications of this theorem include the Fibonacci substitution (since $\operatorname{ind}(\Omega) \geq \operatorname{ind}_\Omega(abaab) \geq 3.2$), the period doubling substitution (since $\operatorname{ind}(\Omega) \geq \operatorname{ind}_\Omega(ab) \geq 3.5$), and many others. Of course, the result does not apply to the Thue–Morse substitution, which is famous mainly because $\operatorname{ind}(\Omega) = 2$. Unfortunately, it is still open whether the point spectrum is almost surely empty in the Thue–Morse case. The Gordon approach fails due to small index, and no other methods are known that yield an almost sure result.[7]

4.5.2. Sturmian and quasi-Sturmian subshifts. Damanik, Killip, and Lenz showed the following result in [58] (see also [59] for a uniform result for almost every Sturmian subshift).

Theorem 4.22. *For every Sturmian subshift Ω, H_ω has empty point spectrum for every $\omega \in \Omega$.*

This result was the culmination of a sequence of partial results. Among those, we single out Sütő [164], who proved empty point spectrum for one α and one ω, Bellissard et al. [14], who proved it for all α and one ω, Delyon–Petritis [73], who proved it for almost all α and almost all ω, and Kaminaga [106], who proved it for all α and almost all ω. Here, $\alpha \in (0,1) \setminus \mathbb{Q}$ denotes the slope associated with a Sturmian subshift. Recall that Sturmian subshifts are in one-to-one correspondence with $(0,1) \setminus \mathbb{Q}$. Here, [14, 58, 59, 164] used Lemma 4.20, whereas [73, 106] used Lemma 4.19.

The extension of Theorem 4.22 to the quasi-Sturmian case was obtained by Damanik and Lenz in [61].

Theorem 4.23. *For every quasi-Sturmian subshift Ω, H_ω has empty point spectrum for every $\omega \in \Omega$.*

[6] See [41] for a precursor dealing with the period doubling case. In this special case it was later shown that the absence of eigenvalues even holds for all $\omega \in \Omega$ [44].

[7] The absence of eigenvalues for a dense G_δ set of ω's can be established in this example and many others using palindromes instead of powers [95]. However, using palindromes one cannot prove the absence of eigenvalues for a full measure set of ω's [72].

4.5.3. Circle maps. Recall that a circle map subshift is determined by the parameters $\alpha \in (0,1) \setminus \mathbb{Q}$ and $\beta \in (0,1)$. It is strictly ergodic and we denote the unique ergodic measure by μ. Delyon and Petritis proved the following in [73].

Theorem 4.24. *For almost every α and every β, the corresponding circle map subshift Ω is such that $\mu(\{\omega \in \Omega : V_\omega$ is a Gordon potential$\}) = 1$. Consequently, $\mu(\{\omega \in \Omega : H_\omega$ has empty point spectrum$\}) = 1$.*

In fact, the full measure set of α values is explicitly described in terms of the continued fraction expansion. The condition was weakened by Kaminaga in [106], still however excluding an explicit zero measure set. This weaker condition was only shown to imply $\mu(\{\omega \in \Omega : V_\omega$ is a Gordon potential$\}) > 0$, which of course is still sufficient to allow one to deduce $\mu(\{\omega \in \Omega : H_\omega$ has empty point spectrum$\}) = 1$.

4.5.4. Interval exchange transformations. Recall that an IET subshift is determined by an irreducible permutation τ and a probability vector λ. Cobo, Gutierrez, and de Oliveira showed the following result in [35] (see also [74]).

Theorem 4.25. *For every irreducible permutation τ and almost every λ, the associated IET subshift Ω is such that $\mu(\{\omega \in \Omega : V_\omega$ is a Gordon potential$\}) = 1$. Consequently, $\mu(\{\omega \in \Omega : H_\omega$ has empty point spectrum$\}) = 1$.*

4.6. Transport properties

Quasicrystal models have behavior that is markedly different from the periodic and random cases in many different respects. In the previous subsections we have seen that the spectrum is typically a zero-measure Cantor set, while for periodic and random potentials it is always given by a finite union of non-degenerate compact intervals. Moreover, the spectral type is typically singular continuous, while it is always absolutely continuous in the periodic case and almost surely pure point in the (one-dimensional) random case.

In this subsection we consider yet another perspective from which the quasicrystal model behavior is expected to differ from the behavior of the periodic and random cases. Namely, we consider the spreading of wave packets under the time-dependent Schrödinger equation. That is, given a Schrödinger operator H_ω and a normalized element ψ of $\ell^2(\mathbb{Z})$, we consider $\psi(t) = e^{-itH_\omega}\psi$, where e^{-itH_ω} is defined via the spectral theorem. Then, $\psi(\cdot)$ satisfies the time-dependent Schrödinger equation $i\partial_t\psi(t) = H_\omega\psi(t)$ with initial condition $\psi(0) = \psi$. The quantum mechanical interpretation is that the probability of finding the quantum particle at site $n \in \mathbb{Z}$ at time $t \in \mathbb{R}$ is given by $a(n,t) := |\langle \delta_n, \psi(t)\rangle|^2$. The initial state is naturally localized in some fixed compact set, up to a small error, since it belongs to $\ell^2(\mathbb{Z})$. More specifically, one is often interested in the initial state $\psi = \delta_0$ (or some δ_n), which is completely localized. After having fixed the initial state, one is then interested in how fast $\psi(t)$ spreads out in space, or more specifically, how long one has to wait until $a(n,t)$ is no longer negligibly small at some distant site n. In general, this is a difficult problem. Questions of this kind are easier to study

for compound quantities; that is, some averaging in n and/or t helps one generate quantities for which interesting statements can be proven.

A popular way to average in time is to consider Cesàro averages,

$$\tilde{a}(n,T) = \frac{1}{T}\int_0^T a(n,t)\,dt = \frac{1}{T}\int_0^T |\langle \delta_n, \psi(t)\rangle|^2\,dt.$$

Let also

$$M_p(t) = \sum_{n\in\mathbb{Z}} (1+|n|^p) a(n,t), \quad \tilde{M}_p(T) = \sum_{n\in\mathbb{Z}} (1+|n|^p)\tilde{a}(n,T), \quad p>0.$$

Notice that for t (resp., T) fixed, $a(\cdot,t)$ and $\tilde{a}(\cdot,T)$ are probability distributions on $\mathbb{Z}$, and hence the quantities above are (1 plus) the pth moment of the respective probability distribution. Here we assume that the initial state is either a Dirac delta function or at least sufficiently well localized so that these moments are finite.

Wave packet spreading then is reflected by growth in time of these moments. To detect power-law growth, one introduces the so-called transport exponents

$$\beta^+(p) = \limsup_{t\to\infty} \frac{\log M_p(t)}{p\log t}, \qquad \beta^-(p) = \liminf_{t\to\infty} \frac{\log M_p(t)}{p\log t},$$

$$\tilde{\beta}^+(p) = \limsup_{T\to\infty} \frac{\log \tilde{M}_p(T)}{p\log T}, \qquad \tilde{\beta}^-(p) = \liminf_{T\to\infty} \frac{\log \tilde{M}_p(T)}{p\log T}.$$

Each of these four functions of p is non-decreasing in p and takes values in the interval $[0,1]$.

In view of the monotonicity of the transport exponents, it is natural to consider their limiting values for small and large values of p. Thus, denote

$$\alpha_\ell^\pm = \lim_{p\downarrow 0} \beta^\pm(p), \qquad \alpha_u^\pm = \lim_{p\uparrow\infty} \beta^\pm(p),$$

$$\tilde{\alpha}_\ell^\pm = \lim_{p\downarrow 0} \tilde{\beta}^\pm(p), \qquad \tilde{\alpha}_u^\pm = \lim_{p\uparrow\infty} \tilde{\beta}^\pm(p).$$

We note that there are other useful ways of capturing wave packet spreading, and refer the reader to [71] for a comprehensive survey.

The transport exponents take the constant value 0 for random potentials and (at least the time-averaged quantities) the constant value 1 for periodic potentials. Thus if one is able to prove the occurrence of fractional values of the transport exponents, one exhibits wave packet spreading that is strictly intermediate between the periodic and random cases. Results of this kind are notoriously difficult to establish. The few known results for quasicrystal models will be described in detail in later sections on Fibonacci and Sturmian potentials.

5. The Fibonacci Hamiltonian

The Fibonacci Hamiltonian is the most prominent model in the study of electronic properties of quasicrystals. It is given by the discrete one-dimensional Schrödinger operator

$$[H_{\lambda,\omega}u](n) = u(n+1) + u(n-1) + \lambda\chi_{[1-\alpha,1)}(n\alpha + \omega \bmod 1)u(n), \tag{5}$$

where $\lambda > 0$ is the coupling constant, $\alpha = \frac{\sqrt{5}-1}{2}$ is the frequency, and $\omega \in [0,1)$ is the phase. An alternative way to obtain the same potential is via the Fibonacci substitution; see Section 4.4.1 above. This operator family has been studied in many papers since the early 1980's (see, e.g., [2, 84, 96, 97, 99, 112, 113, 114, 124, 141, 142, 172] for early works in the physics literature), and numerous fundamental results are known. In this section we describe the current "state of the art" for this model.

5.1. Trace map formalism

Even the earliest papers on the Fibonacci Hamiltonian realized the importance of a certain renormalization procedure in its study, see [112, 141]. This led in particular to the consideration of a certain dynamical system, the so-called trace map, whose properties are closely related to many spectral properties of the operator (5). The existence of the trace map and its connection to spectral properties is a consequence of the invariance of the potential under a substitution rule. This works in great generality; see [3, 43] and references therein.

The one-step transfer matrices associated with the difference equation $H_{\lambda,\omega}u = Eu$ are given by

$$T_{\lambda,\omega}(m,E) = \begin{pmatrix} E - \lambda\chi_{[1-\alpha,1)}(m\alpha + \omega \bmod 1) & -1 \\ 1 & 0 \end{pmatrix}.$$

Denote the Fibonacci numbers by $\{F_k\}$, that is, $F_0 = F_1 = 1$ and $F_{k+1} = F_k + F_{k-1}$ for $k \geq 1$. Then the fact that the potential for zero phase is invariant under the Fibonacci substitution implies that the matrices

$$M_{-1}(E) = \begin{pmatrix} 1 & -\lambda \\ 0 & 1 \end{pmatrix}, \quad M_0(E) = \begin{pmatrix} E & -1 \\ 1 & 0 \end{pmatrix}$$

and

$$M_k(E) = T_{\lambda,0}(F_k, E) \times \cdots \times T_{\lambda,0}(1,E) \quad \text{for } k \geq 1$$

obey the recurrence relations

$$M_{k+1}(E) = M_{k-1}(E)M_k(E)$$

for $k \geq 0$. Passing to the variables

$$x_k(E) = \frac{1}{2}\mathrm{Tr}M_k(E),$$

this in turn implies

$$x_{k+1}(E) = 2x_k(E)x_{k-1}(E) - x_{k-2}(E) \tag{6}$$

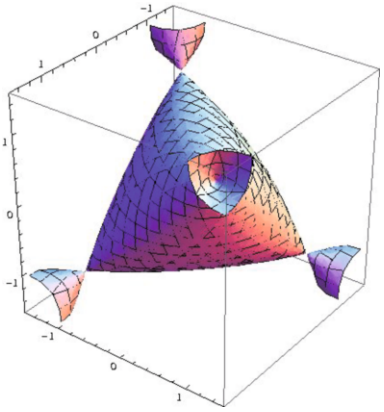

FIGURE 1.
The surface $S_{0.01}$.

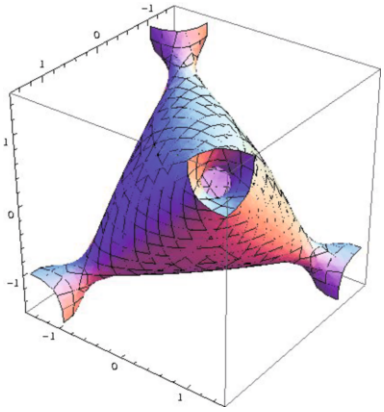

FIGURE 2.
The surface $S_{0.5}$.

for $k \geq 1$, with $x_{-1}(E) = 1$, $x_0(E) = E/2$, and $x_1 = (E - \lambda)/2$. The recursion relation (6) exhibits a conserved quantity; namely, we have

$$x_{k+1}(E)^2 + x_k(E)^2 + x_{k-1}(E)^2 - 2x_{k+1}(E)x_k(E)x_{k-1}(E) - 1 = \frac{\lambda^2}{4} \tag{7}$$

for every $k \geq 0$.

Given these observations, it is then convenient to introduce the *trace map*

$$T : \mathbb{R}^3 \to \mathbb{R}^3, \; T(x, y, z) = (2xy - z, x, y). \tag{8}$$

Aside from the context described here, this map appears in a natural way in problems related to dynamics of mapping classes [85], Fuchsian groups [20], number theory [19], Painlevé sixth equations [31, 103], the Ising model for quasicrystals [15, 91, 175, 176], the Fibonacci quantum walk [154, 155], among others [7, 65, 167, 177]. See [30] or [17] for an algebraic explanation of this universality. We refer the reader also to [98, 156, 157] for further reading on the Fibonacci trace map.

The function

$$G(x, y, z) = x^2 + y^2 + z^2 - 2xyz - 1$$

is invariant under the action of T[8] (which explains (7)), and hence T preserves the family of cubic surfaces[9]

$$S_\lambda = \left\{ (x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 - 2xyz = 1 + \frac{\lambda^2}{4} \right\}. \tag{9}$$

Plots of the surfaces $S_{0.01}$ and $S_{0.5}$ are given in Figures 1 and 2, respectively.

[8]The function $G(x, y, z)$ is usually called the *Fricke character* or *Fricke–Vogt invariant*.
[9]The surface S_0 is called the *Cayley cubic*.

Denote by ℓ_λ the line

$$\ell_\lambda = \left\{ \left(\frac{E-\lambda}{2}, \frac{E}{2}, 1 \right) : E \in \mathbb{R} \right\}.$$

It is easy to check that $\ell_\lambda \subset S_\lambda$.

Sütő proved the following central result in [164].

Theorem 5.1. *An energy E belongs to the spectrum of $H_{\lambda,\omega}$ if and only if the positive semiorbit of the point $\left(\frac{E-\lambda}{2}, \frac{E}{2}, 1\right)$ under iterates of the trace map T is bounded.*

To obtain this theorem, Sütő argued as follows. Denote

$$\sigma_k = \{E \in \mathbb{R} : |x_k(E)| \le 1\}$$

and

$$\Sigma_k = \sigma_k \cup \sigma_{k+1}.$$

These sets depend on the coupling constant λ, and whenever we want to make this dependence explicit, we will write $\sigma_{k,\lambda}$ and $\Sigma_{k,\lambda}$. An analysis of the trace recursion (6) shows that the sets Σ_k are decreasing, and hence it is natural to consider their limit $\tilde{\Sigma} = \bigcap \Sigma_k$. Clearly, if $E \in \tilde{\Sigma}$, then $\{x_n(E)\}$ remains bounded due to (7). On the other hand, the analysis of the trace recursion (6) also yields that whenever $E \notin \Sigma_k$ for some k, then $|x_{n-k}(E)|$ obeys an explicit super-exponentially growing lower bound. That is, the sequence $\{x_n(E)\}$ remains bounded if and only if $E \in \tilde{\Sigma}$. Notice that the point $(\frac{E-\lambda}{2}, \frac{E}{2}, 1)$ is just $(x_1(E), x_0(E), x_{-1}(E))$, so that Theorem 5.1 follows as soon as $\Sigma = \tilde{\Sigma}$ is established. The inclusion $\Sigma \subseteq \tilde{\Sigma}$ holds since σ_k is precisely the spectrum of the canonical periodic approximant of period F_k and the fact that these periodic approximants converge strongly. The inclusion $\tilde{\Sigma} \subseteq \Sigma$ holds since one can use the boundedness of $\{x_n(E)\}$ for $E \in \tilde{\Sigma}$ along with the Gordon lemma to show that no solution for this energy is square-summable at $+\infty$, which implies that E must be in the spectrum.

5.2. Hyperbolicity of the trace map

Let $f : M \to M$ be a diffeomorphism of a Riemannian manifold M. Let us recall that an invariant closed set Λ of the diffeomorphism f is *hyperbolic* if there exists a splitting of a tangent space $T_x M = E^s_x \oplus E^u_x$ at every point $x \in \Lambda$ such that this splitting is invariant under Df, and the differential Df exponentially contracts vectors from stable subspaces $\{E^s_x\}$ and exponentially expands vectors from unstable subspaces $\{E^u_x\}$. A hyperbolic set Λ of a diffeomorphism $f : M \to M$ is *locally maximal* if there exists a neighborhood $U(\Lambda)$ such that

$$\Lambda = \bigcap_{n \in \mathbb{Z}} f^n(U).$$

We will consider diffeomorphisms of a surface, dim $M = 2$, and hyperbolic sets of topological dimension zero. In this case a locally maximal hyperbolic set Λ can be locally represented as a product of "stable" and "unstable" Cantor sets

C^s and C^u. Both Cantor sets C^s and C^u are *dynamically defined*. Dynamically defined Cantor sets have strong self-similar structure and exhibit many nice properties. The formal definition in the general case (when the underlying symbolic system is a general topological Markov chain) is somewhat tedious, and can be found, for example, in [145]. To provide some intuition to the reader, we give here the definition of a dynamically defined Cantor set in the simplest case when the corresponding symbolic dynamical system is a full shift.

Definition 5.2. Let $I \subset \mathbb{R}^1$ be a closed interval. A Cantor set $C \subset I$ is dynamically defined if there are strictly monotone contracting maps $\psi_1, \psi_2, \dots, \psi_k : I \to I, \psi_i(I) \cap \psi_j(I) = \emptyset$ if $i \neq j$, such that $C = \bigcap_{n\in\mathbb{N}} I_n$, where $I_1 = \psi_1(I) \cup \dots \cup \psi_k(I)$ and $I_{n+1} = \psi_1(I_n) \cup \dots \cup \psi_k(I_n)$.

If $\psi_1, \psi_2, \dots, \psi_k$ are $C^{1+\varepsilon}$-functions, then the Cantor set has zero measure, depends continuously on $\psi_1, \dots, \psi_k$, and is "regular" in many other ways. We will be interested in the *Hausdorff dimension* and the *thickness* of the Cantor sets C^s and C^u. Denote the Hausdorff dimension of the set C by $\dim_H C$.

In our case, $\dim_H \Lambda = \dim_H C^s + \dim_H C^u$; see [136, 146]. Moreover, if f depends C^r-smoothly on a parameter, then $\dim_H \Lambda$ is also a smooth function of the parameter; see [134].

Definition 5.3. Let $C \subset \mathbb{R}$ now be an arbitrary Cantor set and denote by I its convex hull. Any connected component of $I \backslash C$ is called a *gap* of C. A *presentation* of C is given by an ordering $\mathcal{U} = \{U_n\}_{n\geq 1}$ of the gaps of C. If $u \in C$ is a boundary point of a gap U of C, we denote by K the connected component of $I \backslash (U_1 \cup U_2 \cup \dots \cup U_n)$ (with n chosen so that $U_n = U$) that contains u and write

$$\tau(C, \mathcal{U}, u) = \frac{|K|}{|U|}.$$

With this notation, the *thickness* $\tau(C)$ and the *denseness* $\theta(C)$ of C are given by

$$\tau(C) = \sup_{\mathcal{U}} \inf_{u} \tau(C, \mathcal{U}, u), \qquad \theta(C) = \inf_{\mathcal{U}} \sup_{u} \tau(C, \mathcal{U}, u). \tag{10}$$

The thickness and the denseness of a Cantor set C are related to the Hausdorff dimension of C by the inequalities (cf. [145, Section 4.2])

$$\frac{\log 2}{\log(2 + \frac{1}{\tau(C)})} \leq \dim_H C \leq \frac{\log 2}{\log(2 + \frac{1}{\theta(C)})}. \tag{11}$$

For more details on thickness, see [77, 138, 145]. An important property of thickness was discovered by Newhouse [140]:

Theorem 5.4. *If C_1 and C_2 are two Cantor sets and $\tau(C_1) \cdot \tau(C_2) \geq 1$, then the sum $C_1 + C_2$ contains an interval. In the special case $C_1 = C_2 =: C$, we have that $\tau(C) \geq 1$ implies that $C + C$ is an interval.*

Consider the restriction $T_\lambda : S_\lambda \to S_\lambda$ of the trace map T from (8) to the invariant surface S_λ, $T_\lambda = T|_{S_\lambda}$. Denote by Ω_λ the set of points in S_λ whose full orbits under T_λ are bounded.

Theorem 5.5. *For every $\lambda > 0$, the set Ω_λ is a locally maximal hyperbolic set of $T_\lambda : S_\lambda \to S_\lambda$. It is homeomorphic to a Cantor set.*

Theorem 5.5 was proved for $\lambda \geq 16$ by Casdagli [33], for small values of λ by Damanik and Gorodetski [50], and finally for all $\lambda > 0$ by Cantat [30].

Since $\ell_\lambda \subset S_\lambda$ the set of points on ℓ_λ whose forward semiorbits are bounded is exactly equal to $\ell_\lambda \cap W^s(\Omega_\lambda)$. Then the spectrum Σ_λ is affine equivalent to the set $\ell_\lambda \cap W^s(\Omega_\lambda)$.

Theorem 5.6. *For every $\lambda > 0$, the line ℓ_λ intersects the leaves of $W^s(\Omega_\lambda)$ transversally.*

This transversality statement was proved for $\lambda \geq 16$ by Casdagli [34, Section 2], and for sufficiently small $\lambda > 0$ by Damanik and Gorodetski [50]. A proof that works for all values of the coupling constant $\lambda > 0$ was given by Damanik, Gorodetski, and Yessen in [57].

Theorem 5.6 allows one to consider the spectrum Σ_λ as a dynamically defined Cantor set. Therefore the following holds.

Corollary 5.7. *For every $\lambda > 0$, the spectrum Σ_λ is a dynamically defined Cantor set, and hence:*

(i) *For every small $\varepsilon > 0$ and every $x \in \sigma(H_{\lambda,\omega})$, we have*

$$\begin{aligned}\dim_H ((x-\varepsilon, x+\varepsilon) \cap \sigma(H_{\lambda,\omega})) &= \dim_B ((x-\varepsilon, x+\varepsilon) \cap \sigma(H_{\lambda,\omega})) \\ &= \dim_H \sigma(H_{\lambda,\omega}) \\ &= \dim_B \sigma(H_{\lambda,\omega}).\end{aligned}$$

(ii) *The Hausdorff dimension $\dim_H \sigma(H_{\lambda,\omega})$ is an analytic function of λ, and is strictly between zero and one.*

5.3. Hausdorff dimension of the spectrum at large coupling

The fact that the box counting dimension of the spectrum exists and coincides with its Hausdorff dimension allows one to determine the asymptotic behavior of this λ-dependent quantity in the large coupling limit. In fact, Damanik, Embree, Gorodetski, and Tcheremchantsev proved the following in [48].

Theorem 5.8. *We have*

$$\lim_{\lambda \to \infty} (\dim \Sigma_\lambda) \cdot \log \lambda = \log(1 + \sqrt{2}).$$

Let us briefly explain how this result is obtained. Recall that the spectrum is related to the spectra of the canonical periodic approximants by

$$\Sigma_\lambda = \bigcap_{k \geq 1} \Sigma_{k,\lambda} = \bigcap_{k \geq 1} \sigma_{k,\lambda} \cup \sigma_{k+1,\lambda}.$$

Since each periodic spectrum $\sigma_{k,\lambda}$ is a finite union of non-degenerate compact intervals and the lengths of these intervals can be shown to be decaying, it is natural to use $\Sigma_{k,\lambda}$ as one possible cover of Σ_λ and estimate the Hausdorff dimension of Σ_λ from above in this way. On the other hand, since each interval of $\sigma_{k,\lambda}$ can be shown to have non-empty intersection with Σ_λ, one can estimate the box counting dimension of Σ_λ from below in this way. We observe how crucial it is that these dimensions coincide here. Thus, the analysis of the participating intervals comes down to proving good estimates for their length.

To estimate the length, one makes use of the following basic fact from one-dimensional Floquet theory. The preimage of the *open* interval $(-1,1)$ under x_k consists of exactly F_k disjoint open intervals, on which x_k is strictly monotone. In fact, in this particular case, the same statement is true for the corresponding *closed* intervals (i.e., the periodic spectra in question have all their gaps open). Thus, the length of one of these intervals (say $I = [a,b]$) can be estimated as follows. Since

$$2 = |x_k(a) - x_k(b)| = \int_a^b |x_k'(E)|\, dE,$$

we have

$$\frac{2}{\max_{E\in I} |x_k'(E)|} \le |I| \le \frac{2}{\min_{E\in I} |x_k'(E)|}.$$

In order to prove estimates for $|x_k'(E)|$, one differentiates the trace recursion (6) and proceeds inductively, making use of the trace invariant (7). This approach was pioneered by Raymond [153] and then used in many subsequent papers. In this inductive approach, it turns out to be important to determine, for a given energy E in one of the intervals of $\sigma_{k,\lambda}$, in how many of the earlier sets $\sigma_{k',\lambda}$, $k' < k$, the energy E in question lies. This gives rise to a combinatorial question that was completely answered in [48]. Combining these combinatorial results with the length estimates one can prove in this way for the intervals in question, the overall strategy above yields the following specific estimates:

$$\dim_H \Sigma_\lambda \le \frac{\log(1+\sqrt{2})}{\log\left(\frac{1}{2}\left[(\lambda-4)+\sqrt{(\lambda-4)^2-12}\right]\right)} \qquad \text{for } \lambda \ge 8, \tag{12}$$

$$\dim_B^- \Sigma_\lambda \ge \frac{\log(1+\sqrt{2})}{\log(2\lambda+22)} \qquad \text{for } \lambda > 4. \tag{13}$$

Theorem 5.8 is then a direct consequence of these estimates and the fact that the Hausdorff dimension and the box counting dimension of Σ_λ are equal.

5.4. Quantitative characteristics of the spectrum at small coupling

Fractal properties of Σ_λ for small λ were studied in [52]. Among many other things, that paper established the following pair of theorems.

Theorem 5.9. *We have*

$$\lim_{\lambda\to 0} \dim \Sigma_\lambda = 1.$$

More precisely, there are constants $C_1, C_2 > 0$ *such that*

$$1 - C_1\lambda \leq \dim \Sigma_\lambda \leq 1 - C_2\lambda$$

for $\lambda > 0$ *sufficiently small.*

Theorem 5.10. *We have*

$$\lim_{\lambda \to 0} \tau(\Sigma_\lambda) = \infty.$$

More precisely, there are constants $C_3, C_4 > 0$ *such that*

$$C_3\lambda^{-1} \leq \tau(\Sigma_\lambda) \leq \theta(\Sigma_\lambda) \leq C_4\lambda^{-1}$$

for $\lambda > 0$ *sufficiently small.*

Theorem 5.9 is a consequence of the connection (11) between the Hausdorff dimension of a Cantor set and its denseness and thickness, along with the estimates for the latter quantities provided by Theorem 5.10.

Let us briefly explain how Theorem 5.10 can be obtained. The Cayley cubic S_0 (cf. (9)) has four conic singularities and can be represented as a union of a two-dimensional sphere (with four conic singularities) and four unbounded components. The restriction of the trace map to the sphere is a pseudo-Anosov map (a factor of a hyperbolic map of a two-torus), and its Markov partition can be presented explicitly (see [33] or [50, 52]). For small values of λ, the map $T : S_\lambda \to S_\lambda$ "inherits" the hyperbolicity of this pseudo-Anosov map everywhere away from the singularities. The dynamics near the singularities must be considered separately. Consider the dynamics of T near one of the singularities, say, near the point $p = (1, 1, 1)$. The set $\mathrm{Per}_2(T)$ of periodic orbits of period two is a smooth curve that contains the point p and intersects S_λ at two points (denote them by $p_1(\lambda)$ and $p_2(\lambda)$) for $\lambda > 0$. Finite pieces of stable and unstable manifolds of $p_1(\lambda)$ and $p_2(\lambda)$ are a distance of order λ from each other. In order to estimate the thickness (and the denseness) of the spectrum Σ_λ, we notice first that the Markov partition for $T : S_0 \to S_0$ can be continuously extended to a Markov partition for $T : S_\lambda \to S_\lambda$. The extended Markov partition is formed by finite parts of the stable and unstable manifolds of $p_1(\lambda)$, $p_2(\lambda)$, and the other six periodic points that are continuations of the three remaining singularities. Therefore the size of the elements of these Markov partitions remains bounded, and the size of the distance between them is of order λ. The natural approach now is to use the distortion property (see, e.g., [145]) to show that for the iterated Markov partition, the ratio of the distance between the elements to the size of an element is of the same order. The main technical problem here is again the dynamics of the trace map near the singularities, since the curvature of S_λ is very large there for small λ. Nevertheless, one can still estimate the distortion that is obtained during a transition through a neighborhood of a singularity and prove boundedness of the distortion for arbitrarily large iterates of the trace map. This implies Theorem 5.10.

5.5. The density of states measure

Let us now turn to the formulation of results involving the integrated density of states, a quantity of fundamental importance associated with an ergodic family of Schrödinger operators. The integrated density of states (IDS) was introduced in Section 3.2 in a more general context, and represents the distribution function of a density of states measure – a measure supported on the spectrum and, in particular, reflecting the asymptotic distribution of eigenvalues of finite-dimensional approximations.

Denote the density of states measure of the Fibonacci Hamiltonian for a given coupling constant λ by dN_λ. Repeating the definition from Section 3.2 in this particular case, we have

$$N_\lambda(E) = \lim_{n\to\infty} \frac{\#\{\text{eigenvalues of } H_{\lambda,\omega,[1,n]} \text{ that are } \le E\}}{n}, \tag{14}$$

where $H_{\lambda,\omega,[1,n]}$ is the restriction of $H_{\lambda,\omega}$ to the interval $[1,n]$ with Dirichlet boundary conditions, and the limit does not actually depend on the phase ω.

It is interesting to analyze the regularity of the density of states measure. This question was studied for general potentials [37, 38, 39, 83, 125], random potentials [29, 162], and analytic quasi-periodic potentials [5, 23, 24, 25, 86, 87, 89, 160]. In the case of Fibonacci Hamiltonian, the IDS is Hölder continuous.

Theorem 5.11. *For every $\lambda > 0$, there exist $C_\lambda < \infty$ and $\gamma_\lambda > 0$ such that*

$$|N_\lambda(E_1) - N_\lambda(E_2)| \le C_\lambda |E_1 - E_2|^{\gamma_\lambda}$$

for every E_1, E_2 with $|E_1 - E_2| < 1$.

This follows directly from [58]; see also [40, 60, 99, 100, 104] for some previous related results.

It is also interesting to obtain the asymptotics of the optimal Hölder exponent for large and small couplings. In the large coupling regime, we have the following [54] (recall that $\alpha = \frac{\sqrt{5}-1}{2}$).

Theorem 5.12.

(a) *Suppose $\lambda > 4$. Then for every*

$$\gamma < \frac{3\log(\alpha^{-1})}{2\log(2\lambda+22)},$$

there is some $\delta > 0$ such that the IDS associated with the family of Fibonacci Hamiltonians satisfies

$$|N_\lambda(E_1) - N_\lambda(E_2)| \le |E_1 - E_2|^{\gamma}$$

for every E_1, E_2 with $|E_1 - E_2| < \delta$.

(b) *Suppose $\lambda \ge 8$. Then for every*

$$\tilde{\gamma} > \frac{3\log(\alpha^{-1})}{2\log\left(\frac{1}{2}\left((\lambda-4)+\sqrt{(\lambda-4)^2-12}\right)\right)}$$

and every $0 < \delta < 1$*, there are* E_1, E_2 *with* $0 < |E_1 - E_2| < \delta$ *such that*

$$|N_\lambda(E_1) - N_\lambda(E_2)| \geq |E_1 - E_2|^{\tilde{\gamma}}.$$

Corollary 5.13. *The optimal Hölder exponent* γ *behaves asymptotically as* $\frac{3\log(\alpha^{-1})}{2\log\lambda}$ *in the large coupling regime.*

The proof is based on the self-similarity of the spectrum and an analysis of the periodic approximants (in the spirit of the proof of Theorem 5.8).

In the small coupling regime, we have the following [54]:

Theorem 5.14. *The integrated density of states* $N_\lambda(\cdot)$ *is Hölder continuous with Hölder exponent* γ_λ*, where* $\gamma_\lambda \to \frac{1}{2}$ *as* $\lambda \to 0$*, and* $\gamma_\lambda < \frac{1}{2}$ *for small* $\lambda > 0$*.*

More precisely:

(a) *For any* $\gamma \in (0, \frac{1}{2})$*, there exists* $\lambda_0 > 0$ *such that for any* $\lambda \in (0, \lambda_0)$*, there exists* $\delta > 0$ *such that*

$$|N_\lambda(E_1) - N_\lambda(E_2)| \leq |E_1 - E_2|^{\gamma}$$

for every E_1, E_2 *with* $|E_1 - E_2| < \delta$*;*

(b) *For any sufficiently small* $\lambda > 0$*, there exists* $\tilde{\gamma} = \tilde{\gamma}(\lambda) < \frac{1}{2}$ *such that for every* $\delta > 0$*, there are* E_1, E_2 *with* $0 < |E_1 - E_2| < \delta$ *and*

$$|N_\lambda(E_1) - N_\lambda(E_2)| \geq |E_1 - E_2|^{\tilde{\gamma}}.$$

The proof uses the trace map formalism and a relation between the IDS of $H_{\lambda,\omega}$ and the measure of maximal entropy for the trace map T_λ. Namely, the density of states measure is proportional to the projection (along the stable manifolds) to ℓ_λ of the normalized restriction of the measure of maximal entropy $\mu_{\max}(T_\lambda)$ to an element of the Markov partition. After that, the proof uses a comparison of expansion rates of T_λ and T_0 (and is reminiscent of the proof of Hölder continuity of conjugacies between two hyperbolic dynamical systems).

Another interesting feature of the Fibonacci Hamiltonian is the uniform scaling of the density of states measure. Namely, the following result (that summarizes the results from [57], [53], and [150]) holds.

Theorem 5.15. *For every* $\lambda > 0$*, there is* $d_\lambda \in (0, 1)$ *so that the density of states measure* dN_λ *is of exact dimension* d_λ*, that is, for* dN_λ*-almost every* $E \in \mathbb{R}$*, we have*

$$\lim_{\varepsilon \downarrow 0} \frac{\log N_\lambda(E - \varepsilon, E + \varepsilon)}{\log \varepsilon} = d_\lambda.$$

Moreover, in $(0, \lambda_0)$*,* d_λ *is an analytic function of* λ*, and*

$$\lim_{\lambda \downarrow 0} d_\lambda = 1.$$

The proof is based on the relation between dN_λ and $\mu_{\max}(T_\lambda)$, and the exact dimensionality of hyperbolic measures [8, 119, 149].

The Hausdorff dimension of the spectrum is an upper bound for d_λ, but a priori it is not clear whether these numbers must coincide. Barry Simon conjectured

that for a large class of models these quantities must be different.[10] The next result by Damanik, Gorodetski, and Yessen [57] shows that this conjecture is true (see also [53] for an earlier partial result).

Theorem 5.16. *For every $\lambda > 0$, we have $d_\lambda < \dim_H \Sigma_\lambda$.*

The proof is based on the comparison of the measure of maximal entropy for T_λ (which is "responsible" for d_λ) and the equilibrium measure for the potential given by minus the log of the expansion rate. The Hausdorff dimension of the unstable projection of the latter is equal to $\dim_H \Sigma_\lambda$, and the thermodynamical description of this measure (see [136]) implies that for any other ergodic invariant measure, the dimension of its unstable projection is strictly smaller. In order to prove that those two measures are actually different, one uses the fact that the measure of maximal entropy is an equilibrium measure that corresponds to zero potential. Therefore it is enough to show that the two potentials under consideration are not cohomological, which can be done using a comparison of multipliers of different periodic orbits of T_λ.

5.6. Gap opening and gap labeling

The spectrum Σ_λ jumps from being an interval for $\lambda = 0$ to being a zero-measure Cantor set for $\lambda > 0$. Hence, as the potential is turned on, a dense set of gaps opens immediately. It is natural to ask about the size of these gaps; see [13]. These gap openings were studied in [10] for the Thue–Morse potential (where the gaps open as a power of λ) and in [12] for the period doubling potential (where some gaps open linearly, and some others are superexponentially small in λ). In the Fibonacci case, all gaps open linearly [52, 57]:

Theorem 5.17. *The boundary points of a gap in the spectrum Σ_λ depend smoothly on the coupling constant λ. Moreover, given any one-parameter continuous family $\{U_\lambda\}_{\lambda>0}$ of gaps of Σ_λ,[11] we have that*

$$\lim_{\lambda \to 0} \frac{|U_\lambda|}{|\lambda|}$$

exists and belongs to $(0, \infty)$.

Theorem 5.17 follows again from dynamical properties of the trace map. Namely, each singularity of the Cayley cubic S_0 gives birth to two periodic points on the surface S_λ, $\lambda > 0$. The distance between the periodic points is of order λ. The stable manifolds of these periodic points "cut" gaps in ℓ_λ that correspond to gaps in the spectrum. The curves formed by the families of the periodic points are normally hyperbolic manifolds of the trace map, and hence (see [92, 152]) their

[10] The conjecture does not appear anywhere in print, but it was popularized by Barry Simon in many talks given by him in the past four years.

[11] By a continuous family $\{U_\lambda\}_{\lambda>0}$ of gaps of Σ_λ we mean that U_λ is a bounded connected component of $\mathbb{R} \setminus \Sigma_\lambda$ and the left endpoint and the right endpoint of U_λ each depend continuously on λ.

strong stable manifolds form a C^1 foliation. This implies that the size of each gap is also of order λ (as $\lambda \to 0$), and Theorem 5.17 follows.

The limit in Theorem 5.17 certainly depends on the family of gaps chosen. In order to study this dependence, one needs to use some labeling of the gaps. As is well known, the density of states produces such a gap labeling. That is, one can identify a canonical set of gap labels, which is only associated with the underlying dynamics (in this case, an irrational rotation of the circle or the shift-transformation on a substitution-generated subshift over two symbols), in such a way that the value of $N(E, \lambda)$ for $E \in \mathbb{R} \setminus \Sigma_\lambda$ must belong to this canonical set. In the Fibonacci case, this set is well known (see, e.g., [13, Eq. (6.7)]) and the general gap labeling theorem specializes to the following statement:

$$\{N(E, \lambda) : E \in \mathbb{R} \setminus \Sigma_\lambda\} \subseteq \{\{m\alpha\} : m \in \mathbb{Z}\} \cup \{1\} \tag{15}$$

for every $\lambda \neq 0$. Here $\{m\alpha\}$ denotes the fractional part of $m\alpha$, that is, $\{m\alpha\} = m\alpha - \lfloor m\alpha \rfloor$. Notice that the set of gap labels is indeed λ-independent and only depends on the value of α from the underlying circle rotation. Since α is irrational, the set of gap labels is dense. In general, a dense set of gap labels is indicative of a Cantor spectrum and hence a common (and attractive) stronger version of proving Cantor spectrum is to show that the operator "has all its gaps open." For example, the Ten Martini Problem for the almost Mathieu operator is to show Cantor spectrum, while the Dry Ten Martini Problem is to show that all labels correspond to gaps in the spectrum. The former problem has been completely solved [4], while the latter has not yet been completely settled. Indeed, it is in general a hard problem to show that all labels given by the gap labeling theorem correspond to gaps, and there are only few results of this kind. It turns out that the stronger (or "dry") form of Cantor spectrum holds for the Fibonacci Hamiltonian [57]:

Theorem 5.18. *For every $\lambda > 0$, all gaps allowed by the gap labeling theorem are open. That is,*

$$\{N(E, \lambda) : E \in \mathbb{R} \setminus \Sigma_\lambda\} = \{\{m\alpha\} : m \in \mathbb{Z}\} \cup \{1\}. \tag{16}$$

Earlier, (16) was shown for $\lambda > 4$ by Raymond [153], and for $\lambda > 0$ sufficiently small by Damanik and Gorodetski [52].

Using the gap labeling, we can refine the statement of Theorem 5.17. For $m \in \mathbb{Z} \setminus \{0\}$, denote by $U_m(\lambda)$ the gap of Σ_λ where the integrated density of states takes the value $\{m\alpha\}$. Then, the following result from [52] holds:

Theorem 5.19. *There is a finite constant C^* such that for every $m \in \mathbb{Z} \setminus \{0\}$,*

$$\lim_{\lambda \to 0} \frac{|U_m(\lambda)|}{|\lambda|} = \frac{C_m}{|m|}$$

for a suitable $C_m \in (0, C^*)$.

To see why Theorem 5.19 holds, notice that each family of gaps converges (as $\lambda \to 0$) to a point of intersection of ℓ_0 with a stable manifold of one of the singularities. The intersections that have larger labels are in a sense "produced"

from intersections with smaller labels by the action of the inverse of the trace map. For gaps with small labels, we know from Theorem 5.17 that $\lim_{\lambda\to 0} \frac{|U_m(\lambda)|}{|\lambda|} < C^*$ for some constant $C^* > 0$. The length (in coordinates on the two-torus covering S_0) of the piece of the stable manifold from the singularity to the point of intersection after k applications of the map is of order $\left(\frac{1+\sqrt{5}}{2}\right)^k \sim |m|$, and the contraction that will be applied to the gap is of order

$$\left(\frac{\sqrt{5}-1}{2}\right)^k \sim \left(\frac{\sqrt{5}-1}{2}\right)^{\frac{\log|m|}{\log\left(\frac{1+\sqrt{5}}{2}\right)}} = \frac{1}{|m|}.$$

5.7. Transport properties

There is a substantial number of papers that investigate the transport exponents associated with the Fibonacci Hamiltonian; see, for example, [16, 40, 45, 48, 58, 66, 67, 68, 69, 70, 104, 110]. While we won't describe all the known results, we want to at least highlight some of them and put them in perspective. As pointed out earlier, one of the fascinating features of quasicrystal models is that the intermediate nature of their aperiodic order between periodic and random is reflected in a number of ways, be it through the spectrum (by spectral measures being purely singular continuous) or through transport behavior. Here we want to address the latter point. All the papers listed above have the goal of proving estimates that show that the transport properties of the Fibonacci Hamiltonian are markedly different from those of periodic or random media.

Since there is ballistic transport (all transport exponents are equal to one) in the periodic case and no transport (all transport exponents are equal to zero) in the random case, one therefore wants to show that the transport exponents take values in the open interval $(0, 1)$. Proving non-trivial lower bounds turns out to be comparatively easier and was accomplished in the late 1990's [40, 104] for zero phase. Several subsequent papers then went on to extend the lower bound to all phases and improved the estimates [48, 58, 66, 67, 68, 70]. Upper bounds for transport exponents, on the other hand, proved to be elusive for some time. Note a key difference here: to bound transport exponents from below, one "only" has to show that some portion of the wave packet moves sufficiently fast. On the other hand, to bound transport exponents from above, one essentially has to control the entire wave packet and show that it does not move too fast (i.e., ballistically). Thus, it is potentially easier to prove upper bounds on transport that are dual to the type of lower bound that had been established, and this indeed turned out to be the case. The papers [45, 110] showed that at least some non-trivial portion of the wave packet moves slowly. Full control and hence genuine upper bounds for transport exponents were finally obtained in 2007 and later [16, 69, 70].

Let us now state some of the transport results explicitly. Some general remarks that should be made are the following:

(a) Almost all results concern time-averaged quantities (i.e., the exponents $\tilde{\beta}^{\pm}(p)$ defined in Section 4.6).
(b) Most papers focus on the case $\psi(0) = \delta_0$. We will limit our attention here to this case as well.
(c) The optimality of the known estimates improves when p and/or λ are large. In particular, the bounds are known to be tight in the limit $\lambda, p \uparrow \infty$.
(d) For finite values of λ and p, the method of choice to obtain the best-known bound varies.
(e) For λ and p large enough, the transport exponent may exceed the dimension of the spectrum.

Here is a result from [68] that establishes the best-known estimates for zero phase and given λ and p:

Theorem 5.20. *Suppose $\lambda > 0$ and set*

$$\gamma = D \log(2 + \sqrt{8 + \lambda^2})$$

(*where D is some universal constant*) *and*

$$\kappa = \log \left[\frac{\sqrt{17}}{20 \log(1 + \alpha)} \right].$$

Then, the time-averaged transport exponent corresponding to the initial state $\psi(0) = \delta_0$ and zero-phase Fibonacci Hamiltonian $H_{\lambda,0}$ obey

$$\tilde{\beta}^{\pm}(p) \geq \begin{cases} \frac{p+2\kappa}{(p+1)(\gamma+\kappa+1/2)}, & p \leq 2\gamma + 1; \\ \frac{1}{\gamma+1}, & p > 2\gamma + 1. \end{cases} \tag{17}$$

Here is a result from [69, 70] that concerns the regime of large λ and p:

Theorem 5.21. *Consider the Fibonacci Hamiltonian $H_{\lambda,\omega}$ and the initial state $\psi(0) = \delta_0$. For $\lambda > \sqrt{24}$, we have*

$$\tilde{\alpha}_u^{\pm} \geq \frac{2 \log(1 + \alpha)}{\log(2\lambda + 22)},$$

and for $\lambda \geq 8$, we have

$$\tilde{\alpha}_u^{\pm} \leq \frac{2 \log(1 + \alpha)}{\log\left(\frac{1}{2}\left[(\lambda - 4) + \sqrt{(\lambda - 4)^2 - 12}\right]\right)}.$$

Both estimates holds uniformly in ω. In particular,

$$\lim_{\lambda \to \infty} \tilde{\alpha}_u^{\pm} \cdot \log \lambda = 2 \log(1 + \alpha),$$

and convergence is uniform in ω.

In fact, the upper bound can be proved also for the non-time-averaged quantities, as shown in [70].

Theorem 5.22. *Consider the Fibonacci Hamiltonian $H_{\lambda,\omega}$ and the initial state $\psi(0) = \delta_0$. For $\lambda \geq 8$ and uniformly in ω, we have*

$$\alpha_u^{\pm} \leq \frac{2\log(1+\alpha)}{\log\left(\frac{1}{2}\left[(\lambda-4)+\sqrt{(\lambda-4)^2-12}\right]\right)}.$$

Some other estimates on transport exponents were obtained recently using different methods in [57].

5.8. Connections between spectral characteristics and dynamical quantities

In [57] explicit relations between spectral quantities for the Fibonacci Hamiltonian and the dynamical characteristics of the Fibonacci trace map were obtained. In the next theorem, $\mu_{\lambda,\max}$ denotes the measure of maximal entropy of $T_\lambda|_{\Omega_\lambda}$ and μ_λ denotes the equilibrium measure of $T_\lambda|_{\Omega_\lambda}$ that corresponds to the potential $-\dim_H \Sigma_\lambda \cdot \log \|DT_\lambda|_{E^u}\|$. Recall that α denotes the inverse of the golden ratio.

Theorem 5.23. *For every $\lambda > 0$, we have*

$$\tilde{\alpha}_u^{\pm}(\lambda) = \frac{\log(1+\alpha)}{\inf_{p\in \mathrm{Per}(T_\lambda)} \mathrm{Lyap}^u(p)}, \tag{18}$$

$$\dim_H \Sigma_\lambda = \frac{h_{\mu_\lambda}}{\mathrm{Lyap}^u \mu_\lambda}, \tag{19}$$

$$\dim_H \nu_\lambda = \dim_H \mu_{\lambda,\max} = \frac{h_{\mathrm{top}}(T_\lambda)}{\mathrm{Lyap}^u \mu_{\lambda,\max}} = \frac{\log(1+\alpha)}{\mathrm{Lyap}^u \mu_{\lambda,\max}}, \tag{20}$$

$$\gamma_\lambda = \frac{\log(1+\alpha)}{\sup_{p\in \mathrm{Per}(T_\lambda)} \mathrm{Lyap}^u(p)}. \tag{21}$$

The following theorem from [57] shows that for the Fibonacci Hamiltonian and every value of the coupling constant, the four quantities satisfy strict inequalities.

Theorem 5.24. *For every $\lambda > 0$, we have*

$$\gamma_\lambda < \dim_H \nu_\lambda < \dim_H \Sigma_\lambda < \tilde{\alpha}_u^{\pm}(\lambda). \tag{22}$$

The particular inequality $\dim_H \nu_\lambda < \dim_H \Sigma_\lambda$ in (22) establishes a conjecture of Barry Simon,[12] which was made based on an analogy with work of Makarov and Volberg [132, 133, 169]; see [53] for a more detailed discussion. The inequality

$$\dim_H \Sigma_\lambda < \tilde{\alpha}_u^{\pm}(\lambda) \tag{23}$$

in (22) is related to a question of Yoram Last. He asked in [118] whether in general $\dim_H \Sigma_\lambda$ bounds $\tilde{\alpha}_u^{\pm}(\lambda)$ from above and conjectured that the answer is no. The inequality (23) confirms this. See [70] and [55] for earlier partial results.

The identities in Theorem 5.23 are instrumental in the proof of Theorem 5.24. Indeed, once the identities (18)–(21) are established, Theorem 5.24 can be proved

[12]The conjecture does not appear anywhere in print, but it was popularized by Barry Simon in many talks given by him in the past four years.

using the thermodynamic formalism, which we will describe next. Define $\phi : \Omega_\lambda \to \mathbb{R}$ by $\phi(x) = -\log \|DT_\lambda(x)|_{E^u}\|$ and consider the pressure function (sometimes called the Bowen function) $P : t \mapsto P(t\phi)$, where $P(\psi)$ is the topological pressure.[13] This function has been heavily studied; the next statement summarizes some known results; compare [28, 109, 143, 158, 170, 171].

Proposition 5.25. *Suppose that $\sigma_A : \Sigma_A \to \Sigma_A$ is a topological Markov chain defined by a transitive 0–1 matrix A, and $\phi : \Sigma_A \to \mathbb{R}$ is a Hölder continuous function. Denote by $\mathfrak{M}$ the space of σ_A-invariant Borel probability measures. Then, the following statements hold.*

(1) *Variational principle: $P(t\phi) = \sup_{\mu\in\mathfrak{M}} \left\{h_\mu + t\int \phi\, d\mu\right\}$.*
(2) *For every $t \in \mathbb{R}$, there exists a unique invariant measure $\mu_t \in \mathfrak{M}$ (the equilibrium state) such that $P(t\phi) = h_{\mu_t} + t\int \phi\, d\mu_t$.*
(3) *$P(t\phi)$ is a real analytic function of t.*
(4) *If ϕ is cohomological to a constant, then $P(t\phi)$ is a linear function; if ϕ is not cohomological to a constant, then $P(t\phi)$ is strictly convex and decreasing.*
(5) *For every $t_0 \in \mathbb{R}$, the line $h_{\mu_{t_0}} + t\int \phi\, d\mu_{t_0}$ is tangent to the graph of the function $P(t\phi)$ at the point $(t_0, P(t_0\phi))$.*
(6) *The following limits exist:*

$$\lim_{t\to\infty} \int \phi\, d\mu_t = \sup_{\mu\in\mathfrak{M}} \int \phi\, d\mu, \qquad \lim_{t\to-\infty} \int \phi\, d\mu_t = \inf_{\mu\in\mathfrak{M}} \int \phi\, d\mu.$$

The graph of the function $t \mapsto P(t\phi)$ lies strictly above each of the lines $t \cdot \sup_{\mu\in\mathfrak{M}} \int \phi\, d\mu$ and $t \cdot \inf_{\mu\in\mathfrak{M}} \int \phi\, d\mu$.

Now let us return to our case where $\sigma_A : \Sigma_A \to \Sigma_A$ is conjugate to $T_\lambda|_{\Omega_\lambda}$ and the potential is given by $\phi(x) = -\log \|DT_\lambda(x)|_{E^u}\|$ (suppressing the conjugacy). In [57] it was shown that this potential is not cohomological to a constant. For any $t \in \mathbb{R}$, consider the tangent line to the graph of $P(t)$ at the point $(t, P(t\phi))$. Since $P(t)$ is decreasing, there exists exactly one point of intersection of the tangent line with the t-axis, at the point $t_0 = -\frac{h_{\mu_t}}{\int \phi\, d\mu} = \frac{h_{\mu_t}}{Lyap^u\, \mu_t} = \dim_H \mu_t$. The last equality here is due to [137]. In particular, $\dim_H \mu_{\max} = \dim_H \nu_\lambda$ is given by the point of intersection of the tangent line to the graph of $P(t)$ at the point $(0, h_{top}(T_\lambda))$ with the t-axis. Also, due to Theorem 5.23 the line $h_{\mathrm{top}}(T_\lambda) + t \cdot \inf_{\mu\in\mathfrak{M}} \int \phi\, d\mu$ intersects the t-axis at the point γ_λ, and the line $h_{\mathrm{top}}(T_\lambda) + t \cdot \sup_{\mu\in\mathfrak{M}} \int \phi\, d\mu$ intersects the t-axis at the point $\tilde{\alpha}_u^{\pm}(\lambda)$. Finally, due to [136], the graph of $P(t)$ intersects the t-axis at the point $\dim_H \Sigma_\lambda$. These observations are illustrated in Figure 3 and explain where the strict inequalities in Theorem 5.24 come from once it is shown that ϕ is not cohomological to a constant.

[13] There are many classical books on the thermodynamical formalism; for example, [28, 158, 172]. We also refer the reader to the recent introductory texts [9, 101, 159].

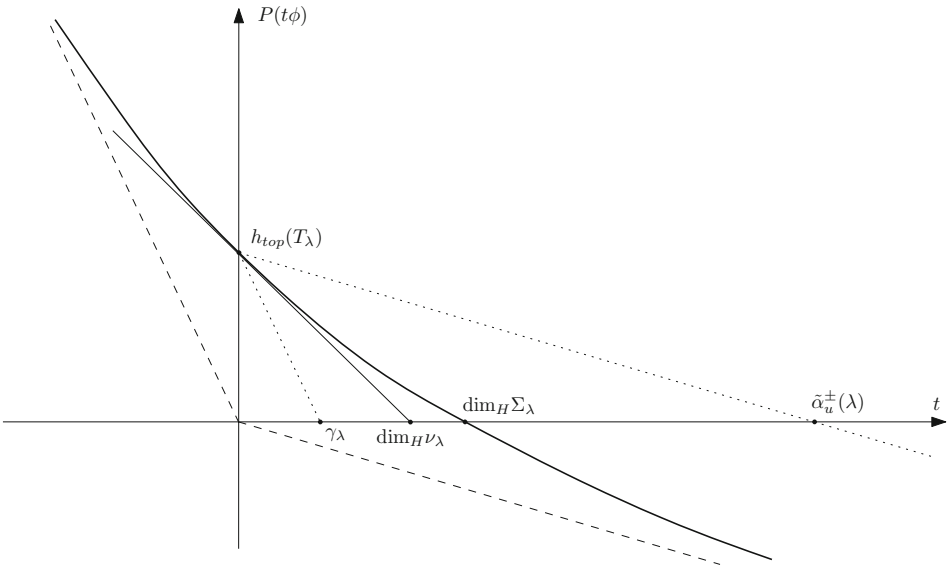

FIGURE 3. Pressure function and spectral characteristics of the Fibonacci Hamiltonian.

5.9. Square and cubic Fibonacci Hamiltonians

Since spectral questions for Schrödinger operators in two (and higher) dimensions are hard to study, it is natural to consider a model where known one-dimensional results can be used. In particular, let us consider the Schrödinger operator

$$\begin{aligned}[H^{(2)}_{\lambda_1,\lambda_2,\omega_1,\omega_2}\psi](m,n) \\ = \psi(m+1,n) + \psi(m-1,n) + \psi(m,n+1) + \psi(m,n-1) \\ + \big(\lambda_1 \chi_{[1-\alpha,1)}(m\alpha+\omega_1 \mod 1) + \lambda_2 \chi_{[1-\alpha,1)}(n\alpha+\omega_2 \mod 1)\big)\,\psi(m,n)\end{aligned} \tag{24}$$

in $\ell^2(\mathbb{Z}^2)$. The theory of tensor products of Hilbert spaces and operators then implies that $\sigma(H^{(2)}_{\lambda_1,\lambda_2,\omega_1,\omega_2}) = \Sigma_{\lambda_1} + \Sigma_{\lambda_2}$ for all ω_1, ω_2. This operator and its spectrum have been studied numerically and heuristically by Even-Dar Mandel and Lifshitz in a series of papers [79, 80, 81] (a similar model was studied by Sire in [163]). Their study suggested that at small coupling, $\Sigma_{\lambda_1} + \Sigma_{\lambda_2}$ is not a Cantor set; quite on the contrary, it has no gaps at all.

It turns out that this is indeed the case [52]:

Theorem 5.26. *For $\lambda_1, \lambda_2 > 0$ sufficiently small, $\sigma(H^{(2)}_{\lambda_1,\lambda_2,\omega_1,\omega_2}) = \Sigma_{\lambda_1} + \Sigma_{\lambda_2}$ is an interval.*

This result follows from the estimates for the thickness of Σ_λ from Theorem 5.10 and Newhouse's Gap Lemma (Theorem 5.4).

Theorem 5.26 should be contrasted with the following result, which is an immediate consequence of Corollary 5.7 and Theorem 5.8.

Theorem 5.27. *For $\lambda_1, \lambda_2 > 0$ sufficiently large, $\sigma(H^{(2)}_{\lambda_1,\lambda_2,\omega_1,\omega_2}) = \Sigma_{\lambda_1} + \Sigma_{\lambda_2}$ is a Cantor set.*

The same statements hold for the cubic Fibonacci Hamiltonian (i.e., the analogously defined Schrödinger operator in $\ell^2(\mathbb{Z}^3)$ with spectrum $\Sigma_{\lambda_1}+\Sigma_{\lambda_2}+\Sigma_{\lambda_3}$). Section 7.3 shows numerical illustrations of the finite approximations $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ and $\Sigma_{k,\lambda}+\Sigma_{k,\lambda}+\Sigma_{k,\lambda}$, along with an exploration of the number of disjoint intervals that make up these sets.

Moreover, the density of states measure of the family $\{H^{(2)}_{\lambda_1,\lambda_2,\omega_1,\omega_2}\}_{\lambda_j\in\mathbb{R},\omega_j\in\mathbb{T}}$ can be expressed as the convolution of the density of states measures associated with the families $\{H_{\lambda_1,\omega_1}\}_{\omega_1\in\mathbb{T}}$ and $\{H_{\lambda_2,\omega_2}\}_{\omega_2\in\mathbb{T}}$, that is,

$$\nu^{(2)}_{\lambda_1,\lambda_2} = \nu_{\lambda_1} * \nu_{\lambda_2}. \tag{25}$$

See the appendix in [56] for further background on separable potentials and operators. The following result was obtained by Damanik, Gorodetski and Solomyak in [56].

Theorem 5.28. *Let $\nu^{(2)}_{\lambda_1,\lambda_2}$ be the density of states measure for the Square Fibonacci Hamiltonian* (24) *with coupling constants λ_1, λ_2. There is $\lambda^* > 0$ such that for almost every pair $(\lambda_1, \lambda_2) \in [0, \lambda^*) \times [0, \lambda^*)$, the measure $\nu^{(2)}_{\lambda_1,\lambda_2}$ is absolutely continuous with respect to Lebesgue measure.*

In fact, it follows from the proof that (with a uniform smallness condition) for every $\lambda_1 \in [0, \lambda^*)$, the measure $\nu^{(2)}_{\lambda_1,\lambda_2}$ is absolutely continuous with respect to the Lebesgue measure for almost every $\lambda_2 \in [0, \lambda^*)$.

6. Sturmian potentials

The Fibonacci potential is a special case of a Sturmian potential. The latter are obtained if α in the definition of the potential, $V(n) = \lambda\chi_{[1-\alpha,1)}(n\alpha+\omega \bmod 1)$, is a general irrational number in $(0, 1)$. The Fibonacci case corresponds to the choice $\alpha = \frac{\sqrt{5}-1}{2}$.

Given an irrational $\alpha \in (0, 1)$, consider its continued fraction expansion

$$\alpha = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cdots}}}$$

with uniquely determined $a_k \in \mathbb{Z}_+$. Truncating the continued fraction expansion of α after k steps yields the rational number p_k/q_k, which is the best rational

approximant of α with denominator bounded by $q_{k+1}-1$. The following recursions hold:

$$p_{k+1} = a_{k+1}p_k + p_{k-1}, \quad p_0 = 0,\ p_1 = 1,$$
$$q_{k+1} = a_{k+1}q_k + q_{k-1}, \quad q_0 = 1,\ q_1 = a_1.$$

(In the Fibonacci case $\alpha = \frac{\sqrt{5}-1}{2}$, we have $a_k \equiv 1$ and $p_k/q_k = F_{k-1}/F_k$.) A number of the results for the Fibonacci Hamiltonian described in the previous section have been generalized to the Sturmian case under suitable assumptions on the continued fraction coefficients $\{a_k\}$. In this section, we explain what these results are, and how the proofs had to be modified.

6.1. Extension of the trace map formalism

Let us the denote the discrete Schrödinger operator on $\ell^2(\mathbb{Z})$ with potential $V(n) = \lambda\chi_{[1-\alpha,1)}(n\alpha+\omega \bmod 1)$ by $H_{\lambda,\alpha,\omega}$. Strong approximation again shows that the spectrum of $H_{\lambda,\alpha,\omega}$ does not depend on ω, and may therefore be denoted by $\Sigma_{\lambda,\alpha}$. The one-step transfer matrices associated with the difference equation $H_{\lambda,\alpha,\omega}u = Eu$ are given by

$$T_{\lambda,\alpha,\omega}(m,E) = \begin{pmatrix} E - \lambda\chi_{[1-\alpha,1)}(m\alpha+\omega \bmod 1) & -1 \\ 1 & 0 \end{pmatrix}.$$

The matrices

$$M_{-1}(E) = \begin{pmatrix} 1 & -\lambda \\ 0 & 1 \end{pmatrix}, \quad M_0(E) = \begin{pmatrix} E & -1 \\ 1 & 0 \end{pmatrix},$$

and

$$M_k(E) = T_{\lambda,\alpha,0}(q_k,E) \times \cdots \times T_{\lambda,\alpha,0}(1,E) \quad \text{for } k \geq 1$$

obey the recurrence relations

$$M_{k+1}(E) = M_{k-1}(E)M_k(E)^{a_{k+1}}$$

for $k \geq 0$; see [14, Proposition 1]. Passing to the variables

$$x_k(E) = \frac{1}{2}\mathrm{Tr}M_k(E),$$

this in turn implies via the Cayley-Hamilton theorem that $x_{k+1}(E)$ can be expressed as an explicit function of (suitable Chebyshev polynomials applied to) $x_k(E), x_{k-1}(E), x_{k-2}(E)$ for $k \geq 1$; see [14, Proposition 2]. These recursion relations exhibit the same conserved quantity as before; namely, with

$$\tilde{x}_{k+1}(E) = \frac{1}{2}\mathrm{Tr}(M_k(E)M_{k-1}(E)),$$

we have

$$\tilde{x}_{k+1}(E)^2 + x_k(E)^2 + x_{k-1}(E)^2 - 2\tilde{x}_{k+1}(E)x_k(E)x_{k-1}(E) - 1 = \frac{\lambda^2}{4}$$

for every $k \geq 0$; see [14, Proposition 3].

6.2. Results obtained via an analysis of the trace recursions

Notice that the key difference with the Fibonacci case is that, in general, the sequence of traces may not be obtained by iterating a single map. In this sense, there is in general no direct analog of the trace map. However, as we have just seen, the underlying structure of recurrence relations extends nicely. The substitute for the dynamical analysis of the Fibonacci trace map will have to lie in studying the dynamics of an initial point under the successive application of a sequence of maps, the elements of which are dictated by the continued fraction expansion of α. These developments are still in their early stages. In the following we will concentrate on the known results that can be established by simply exploiting the recurrence relations, without employing sophisticated tools from dynamical systems theory.

The first result that establishes a clean analogy with the Fibonacci case is the following analog of Theorem 5.1, which was established in [14].

Theorem 6.1. *Fix $\lambda > 0$ and $\alpha \in (0,1)$ irrational. An energy E belongs to the spectrum $\Sigma_{\lambda,\alpha}$ if and only if the sequence $\{x_k(E)\}$ is bounded.*

The proof of Theorem 6.1 follows the same line of reasoning as the proof of Theorem 5.1, which was outlined in the previous section. In particular, one obtains canonical covers of the spectrum, which are useful in the estimation of its dimension. Let us make this explicit. As before, define the sets

$$\sigma_{\lambda,\alpha,k} = \{E \in \mathbb{R} : |x_k(E)| \leq 1\}$$

and

$$\Sigma_{\lambda,\alpha,k} = \sigma_k \cup \sigma_{k+1}.$$

The same reasoning shows that the sets $\Sigma_{\lambda,\alpha,k}$ are decreasing in k and the spectrum is the limiting set, that is,

$$\Sigma_{\lambda,\alpha} = \bigcap_{k \geq 1} \Sigma_{\lambda,\alpha,k};$$

see [14, Proposition 4].

A refinement of this description of the spectrum in the Sturmian case due to Raymond [153] allowed Liu and Wen to obtain the following estimates for the Hausdorff dimension of the spectrum in the large coupling regime [130].

Theorem 6.2. *Suppose $\lambda > 20$ and $\alpha \in (0,1)$ is irrational with continued fraction coefficients $\{a_k\}$. Denote*

$$M_* = \liminf_{k \to \infty} (a_1 \cdots a_k)^{1/k} \in [1, \infty].$$

(a) *If $M_* = \infty$, then $\dim_H \Sigma_{\lambda,\alpha} = 1$.*
(b) *If $M_* < \infty$, then $\dim_H \Sigma_{\lambda,\alpha}$ belongs to the open interval $(0,1)$ and obeys the estimates*

$$\dim_H \Sigma_{\lambda,\alpha} \leq \frac{2 \log M_* + \log 3}{2 \log M_* - \log \frac{3}{\lambda - 8}}$$

and

$$\dim_H \Sigma_{\lambda,\alpha} \geq \max\left\{ \frac{\log 2}{10\log 2 - 3\log\frac{1}{4(\lambda-8)}}, \frac{\log M_* - \log 3}{\log M_* - \log\frac{1}{12(\lambda-8)}} \right\}.$$

A study of the box counting dimension of $\Sigma_{\lambda,\alpha}$ in the case of bounded $\{a_k\}$ was carried out in the follow-up paper [82] by Fan, Liu, and Wen. Among other things, they showed that for $\lambda > 20$, the Hausdorff dimension and the box counting dimension of $\Sigma_{\lambda,\alpha}$ coincide whenever the sequence $\{a_k\}$ is eventually periodic. The analysis of the case of unbounded $\{a_k\}$ was carried out by Liu, Qu, and Wen in [128]. On the one hand, these papers establish the following companion result to Theorem 6.2.

Theorem 6.3. *Suppose $\lambda \geq 24$ and $\alpha \in (0,1)$ is irrational with continued fraction coefficients $\{a_k\}$. Denote*

$$M^* = \limsup_{k\to\infty} (a_1 \cdots a_k)^{1/k} \in [1,\infty].$$

(a) *If $M^* = \infty$, then $\dim_B^+ \Sigma_{\lambda,\alpha} = 1$.*
(b) *If $M^* < \infty$, then $\dim_B^+ \Sigma_{\lambda,\alpha}$ belongs to the open interval $(0,1)$.*

Here $\dim_B^+ S$ denotes the upper box counting dimension of the set S. Note that Theorems 6.2 and 6.3 imply in particular that for suitable choices of α and λ, we may have $\dim_H \Sigma_{\lambda,\alpha} < 1$ and $\dim_B^+ \Sigma_{\lambda,\alpha} = 1$.

On the other hand, Liu, Qu, and Wen also study in [128] the large coupling asymptotics of these dimensions. Namely they show that the limits

$$\lim_{\lambda\to\infty} \dim_H \Sigma_{\lambda,\alpha} \cdot \log\lambda \qquad \text{and} \qquad \lim_{\lambda\to\infty} \dim_B^+ \Sigma_{\lambda,\alpha} \cdot \log\lambda$$

exist, and provide a description of these limits.

The transport exponents in the Sturmian case were studied in the papers [40, 58, 66, 68, 135]. The following result from [68] gives dynamical lower bounds for all values of λ and p, provided α has bounded continued fraction coefficients.

Theorem 6.4. *Suppose $\lambda > 0$ and $\alpha \in (0,1)$ is irrational with $a_k \leq C$. With*

$$\gamma = D\,\log(2+\sqrt{8+\lambda^2}) \cdot \limsup_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} a_k$$

(where D is some universal constant) and

$$\kappa = \frac{\log(\sqrt{17}/4)}{(C+1)^5},$$

the transport exponents associated with the operator $H_{\lambda,\alpha,0}$ and the initial state $\psi(0) = \delta_0$ obey

$$\tilde{\beta}^-(p) \geq \begin{cases} \frac{p+2\kappa}{(p+1)(\gamma+\kappa+1/2)}, & p \leq 2\alpha+1; \\ \frac{1}{\gamma+1}, & p > 2\alpha+1. \end{cases}$$

The following result from [135] gives dynamical upper bounds in the large coupling regime.

Theorem 6.5. *Suppose $\lambda > 20$ and $\alpha \in (0,1)$ is irrational with continued fraction coefficients $\{a_k\}$ and corresponding rational approximants $\{p_k/q_k\}$. Denote*

$$D = \limsup_{k\to\infty} \frac{1}{k} \log q_k.$$

Then, the transport exponents associated with the operator $H_{\lambda,\alpha,0}$ and the initial state $\psi(0) = \delta_0$ obey

$$\tilde{\alpha}_u^{\pm} \leq \frac{2D}{\log \frac{\lambda-8}{3}}.$$

Moreover, if $a_k \geq 2$ for all k, then

$$\tilde{\alpha}_u^{\pm} \leq \frac{D}{\log \frac{\lambda-8}{3}}.$$

7. Numerical results and computational issues

In this section, we provide numerical illustrations of a number of the results described in this survey. These calculations focus on the Fibonacci Hamiltonian, though many could readily be adapted to the Sturmian potentials described in the last section. We begin by studying approximations to the spectrum for the Fibonacci model in one dimension, then investigate estimates of the integrated density of states based on spectra of finite sections of the operator. Finally, we address upper bounds on the spectrum in two and three dimensions. In all cases, we set the phase ω to zero.

7.1. Spectral approximations for the Fibonacci Hamiltonian

We begin by calculating the spectrum σ_k for the kth periodic approximations to the Fibonacci potential. The analysis described in Section 5 suggests several ways to compute σ_k, which turn out to have varying degrees of utility.

Given a candidate energy E, one can test if $E \in \sigma_k$ by iterating the trace recurrence (6) and testing if $|x_k(E)| \leq 1$. In principle, this simple approach enables investigation for arbitrarily large values of k. However, two key obstacles restrict the utility of this method: (i) it does not readily yield the entire set σ_k; (ii) as k increases, the intervals that comprise σ_k become exponentially narrow, beyond the resolution of the standard floating point number system in which such calculations are typically performed. However, this approach can yield some useful results, particularly in the small coupling regime where the decay of the interval widths is most gradual, or when one is only interested in some narrow set of energy values. (This method of calculation was used to produce illustrations in [52].)

To obtain the entire set σ_k, one might instead use the recurrence (6) to construct the degree-F_k polynomial $x_k(E)$, then determine the regions where $|x_k(E)| \leq 1$ by finding the zeros of the polynomials $x_k(E) + 1$ and $x_k(E) - 1$

using a standard root-finding algorithm. For all but the smallest k this approach is untenable. Coefficients of $x_k(E)$ grow exponentially in k; e.g., for $\lambda = 4$,

$$x_6(E) = \tfrac{1}{2}E^{13} - 16E^{12} + \tfrac{435}{2}E^{11} - 1616E^{10} + \tfrac{13905}{2}E^9 - 16272E^8 + 13330E^7 + 20160E^6 - 37133E^5 - 17056E^4 + \tfrac{61013}{2}E^3 + 25104E^2 + \tfrac{13021}{2}E + 560.$$

The magnitude of these coefficients, compounded by the proximity of the roots for larger values of λ and k, leads to inaccurate root calculations, a phenomenon well studied by numerical analysts; see, e.g., [139, 173]. Indeed, it is not uncommon for the computed roots to be so inaccurate as to have significant spurious imaginary parts.

There is a more robust approach to computing the approximate Fibonacci spectrum σ_k. One can view σ_k as the exact spectrum of a related Schrödinger operator with a potential having period F_k. The spectrum of this operator is the union of F_k non-degenerate intervals whose endpoints are given by the eigenvalues of the two F_k-dimensional matrices J_{k+} and J_{k-}:

$$J_{k\pm} = \begin{pmatrix} v_{1,k} & 1 & & & \pm 1 \\ 1 & v_{2,k} & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & v_{F_k-1,k} & 1 \\ \pm 1 & & & 1 & v_{F_k,k} \end{pmatrix},$$

with unspecified entries beyond the tridiagonal section equal to zero; see, e.g., [168, Ch. 7]. Here the potential values $v_{n,k}$ are given by

$$v_{n,k} = \lambda \chi_{[1-F_{k-1}/F_k,1)}(nF_{k-1}/F_k \bmod 1). \tag{26}$$

This approach, which we use for the computations described below, has also been employed in the context of Fibonacci computations by Even-Dar Mandel and Lifshitz [79], and for the almost Mathieu operator by Lamoureux [117].

The standard procedure for computing all the eigenvalues of a symmetric matrix begins by applying a unitary similarity transformation to reduce the matrix to symmetric tridiagonal form.[14] Floating-point arithmetic introduces errors into this process, resulting in the exact tridiagonal reduction of a matrix that differs from the intended matrix by a factor that scales with the precision of the floating point arithmetic system, the coupling constant λ, and the dimension F_k. The eigenvalues of this tridiagonal matrix are then approximated to high accuracy via a procedure known as QR iteration [147]. Remarkably, this iteration does not introduce significant errors beyond those incurred by the reduction to tridiagonal form; for a discussion of this accuracy, see [1, 174]. Overall, this process requires

[14] Methods such as the Lanczos algorithm excel at computing *a few* eigenvalues of large symmetric matrices [147, Ch. 13]. These methods are not feasible here, for *all* eigenvalues of $J_{k\pm}$ are required. However, if one is only interested in a narrow band of energies, these methods can be highly effective.

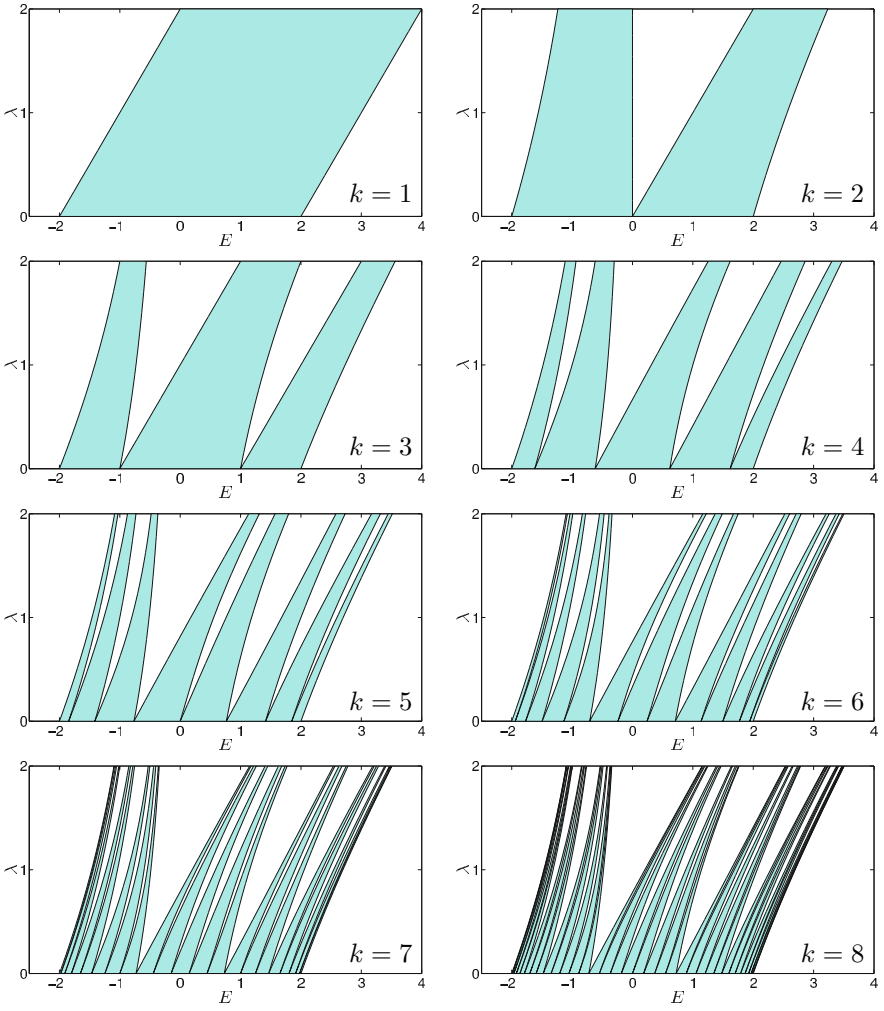

FIGURE 4. Spectra of the periodic approximations $\sigma_{k,\lambda}$ for the Fibonacci Hamiltonian, as a function of $\lambda \in [0, 2]$. For $k = 8$ and all $\lambda > 0$, $\sigma_{k,\lambda}$ is the union of $F_8 = 34$ disjoint intervals.

$O(F_k^3)$ floating point arithmetic operations and the storage of $O(F_k^2)$ floating point numbers. (The conventional procedure for reducing the matrix to tridiagonal form destroys the zero structure present in $J_{k\pm}$.) Of course, the upper estimate $\Sigma_{k,\lambda} = \sigma_{k,\lambda} \cup \sigma_{k+1,\lambda}$ then requires computation of all eigenvalues of four matrices.

Significant insight can be gleaned from numerical calculations involving small to moderate values of k. For example, Figure 4 shows $\sigma_{k,\lambda}$ for $\lambda \in [0, 2]$ and $k = 1, \ldots, 8$, while Figure 5 shows the upper bounds $\Sigma_{k,\lambda}$ for the same range of λ

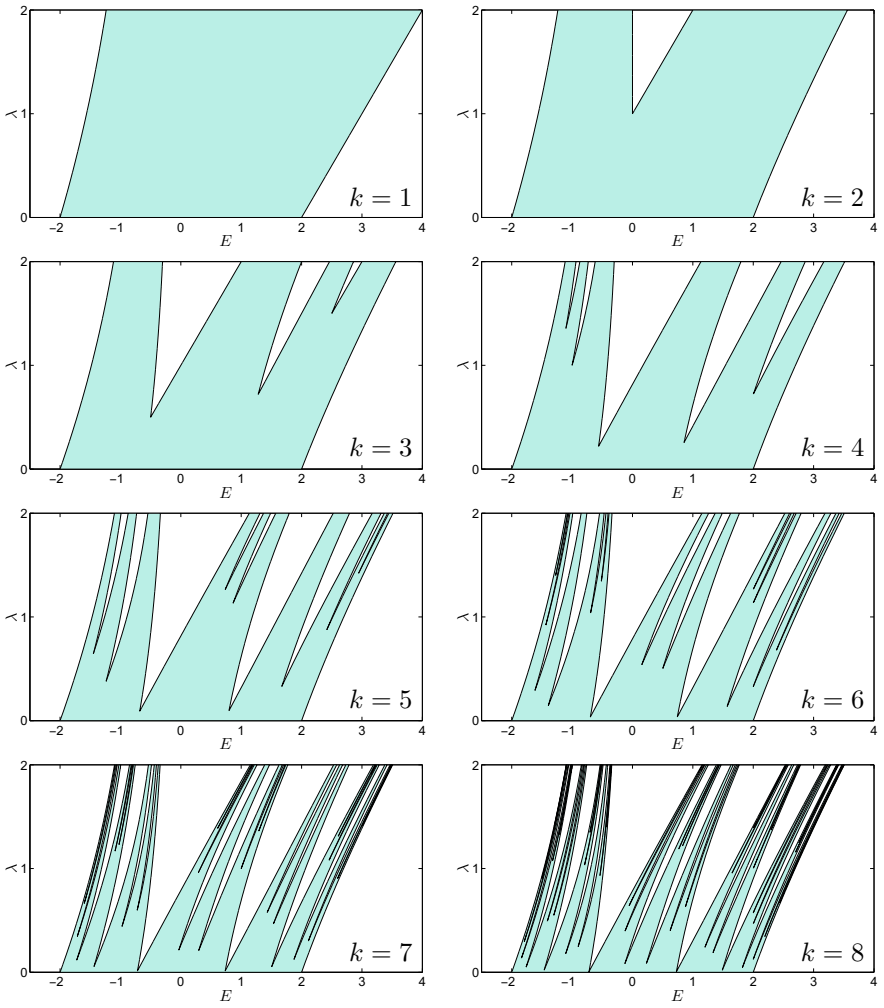

FIGURE 5. The upper bound $\Sigma_{k,\lambda} = \sigma_{k,\lambda} \cup \sigma_{k+1,\lambda}$ on the Fibonacci spectrum, as a function of $\lambda \in [0, 2]$. For $k = 8$ and $\lambda = 2$, $\Sigma_{k,\lambda}$ is the union of 42 disjoint intervals.

and k. Since $\Sigma_{8,\lambda} = \sigma_{8,\lambda} \cup \sigma_{9,\lambda}$, for $\lambda > 0$ the spectrum is the union of 34 and 55 intervals.

To develop conjectures (e.g., regarding dim Σ_λ), one would like to use approximations to Σ_λ for larger k. Two fundamental challenges arise: (i) the $O(F_k^3)$ work and $O(F_k^2)$ storage becomes prohibitive; (ii) while non-degenerate, the intervals in $\sigma_{k,\lambda}$ become exponentially small and exponentially close together. This phenomenon is illustrated in Figure 6. The utility of the numerical results degrade

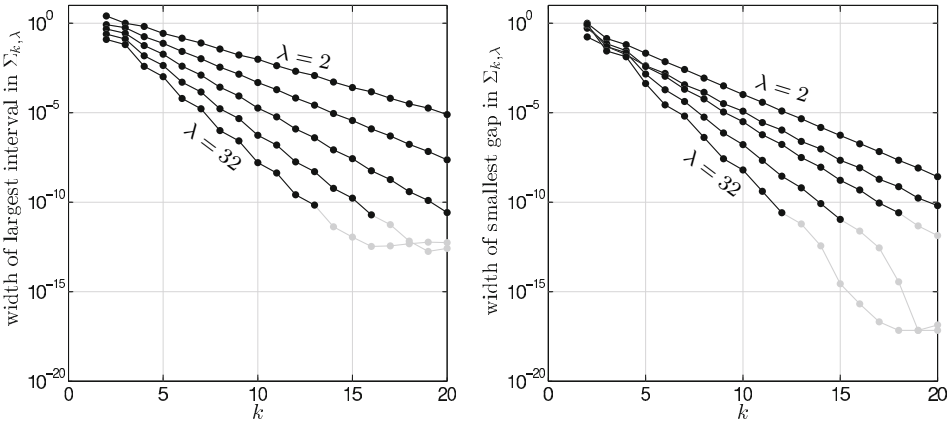

FIGURE 6. Exponential decay of the largest intervals and smallest gaps in the approximations $\Sigma_{k,\lambda}$, for coupling constants $\lambda = 2, 4, 8, 16, 32$, as computed in MATLAB's double-precision floating point arithmetic. The data points that are plotted in gray are likely dominated by computational errors.

when the size of these bands and gaps approaches the order of the error in the numerical computation.[15] On contemporary commodity computers, computation of $\Sigma_{k,\lambda}$ up to roughly $k = 20$ (requiring all eigenvalues of matrices of dimension $F_{20} = 10{,}946$ and $F_{21} = 17{,}711$) is feasible, provided λ is sufficiently small for the results to be accurate. Recently Puelz has proposed an improved approach that ameliorates challenge (i) above by reducing the required work to $O(F_k^2)$ and storage to $O(F_k)$, and challenge (ii) by using extended precision arithmetic [151].

To estimate the box-counting dimension of Σ_λ (assuming it exists), we use the definition

$$\dim_B(S) = \lim_{\varepsilon \to 0} \frac{\log C_S(\varepsilon)}{\log 1/\varepsilon},$$

where $C_S(\varepsilon)$ counts the number of intervals of width ε that intersect S,

$$C_S(\varepsilon) := \#\{j \in \mathbb{Z} : [j\varepsilon, (j+1)\varepsilon) \cap S \neq \emptyset\}.$$

Note that $\dim_B(\Sigma_{k,\lambda}) = 1$ for all k, since $\Sigma_{k,\lambda}$ is the union of finitely many closed intervals. Still, one gains insight into $\dim_B(\Sigma_\lambda)$ from $\log(C_{\Sigma_{k,\lambda}}(\varepsilon))/\log(1/\varepsilon)$ for finite values of ε and various k, as can be seen in Figure 7. For fixed λ, the resulting estimates of $\dim_B(\Sigma_\lambda)$ (taken, e.g., as $\inf_{\varepsilon \in (0,1)} \log(C_{\Sigma_{k,\lambda}}(\varepsilon))/\log(1/\varepsilon)$) apparently improve as k increases; lower values of k are suitable for larger values of λ. However, with this approach it is difficult to accurately estimate the critical value at which

[15] More subtly, the formula (26) incurs significant rounding errors for large n and k, resulting in errors on the diagonal of $J_{k\pm}$ of size λ. For greater accuracy, one should use the equivalent formulation $v_{n,k} = \lambda \chi_{[F_k - F_{k-1}, F_k)}(nF_{k-1} \bmod F_k)$, which is more robust.

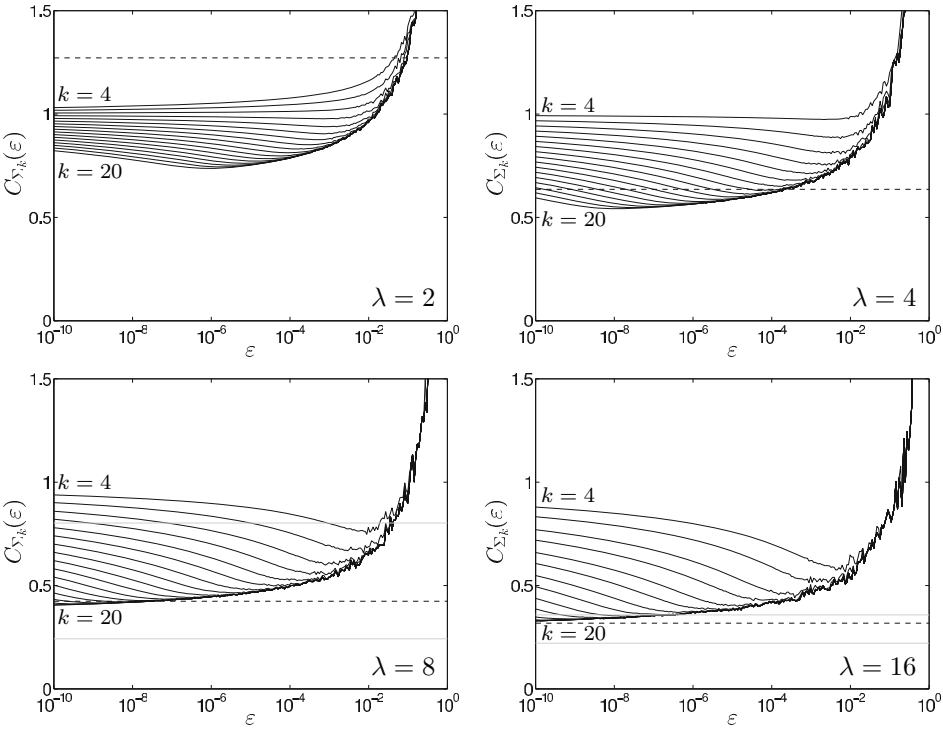

FIGURE 7. Estimates of $\dim_B(\Sigma_\lambda)$ for various values of λ, based on the upper bounds $\Sigma_{k,\lambda}$ for various k. The dashed horizontal line denotes $\log(1+\sqrt{2})/\log(\lambda)$, to which $\dim_B(\Sigma_\lambda)$ tends as $\lambda \to \infty$ (Theorem 5.8). The gray horizontal lines in the bottom plots show the upper and lower bounds (12)–(13).

$\dim_B(\Sigma_\lambda) = 1/2$.[16] (A rough estimate, suggested from Figure 7, is $\lambda \approx 4$; see the discussion preceding Problem 8.6 below.) More accurate approximations will require computations with larger values of k than are feasible with the method described above.

Finally, Figure 8 explores numerical computations of the thickness, defined in (10). As established in Theorem 5.10, the thickness $\tau(\Sigma_\lambda)$ behaves like $1/\lambda$ as $\lambda \downarrow 0$. As λ decreases we see this behavior mirrored in the upper bounds $\Sigma_{k,\lambda}$, up to some point where $\tau(\Sigma_{k,\lambda})$ rapidly increases: $\Sigma_{k,\lambda}$ is the union of no more than $F_k + F_{k+1}$ intervals separated by gaps that diminish as $\lambda \downarrow 0$.

[16] We are interested in this critical value because as soon as $\dim_B(\Sigma_\lambda)$ falls below $1/2$, we can be sure that the sum set $\Sigma_\lambda + \Sigma_\lambda$ is a zero-measure Cantor set, and this is an issue of interest for reasons we will discuss in Subsection 7.3.

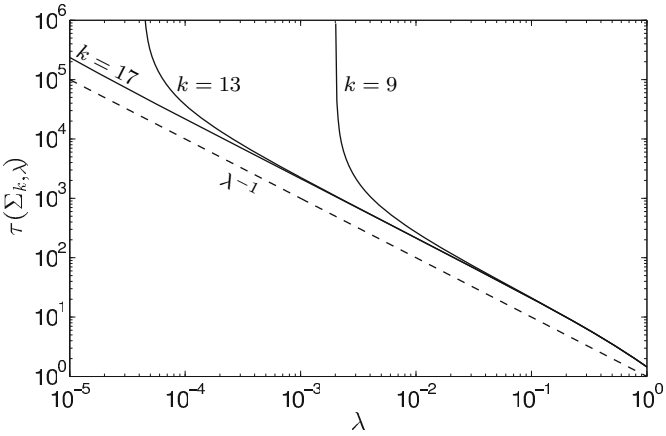

FIGURE 8. Thickness of $\Sigma_{k,\lambda}$ as a function of λ for three values of k, consistent with Theorem 5.10.

7.2. Density of states for the Fibonacci model

We next turn to an investigation of the exponent of Hölder continuity of the integrated density of states (IDS) for the Fibonacci model, discussed in Section 5.5. To estimate $N_\lambda(E)$ in equation (14), one must compute all the eigenvalues of $H_{\lambda,[1,n]}$, the restriction of H_λ to sites $[1, n]$ with Dirichlet boundary conditions. This restriction is an $n \times n$ tridiagonal matrix; because this matrix lacks the corner entries present in $J_{k\pm}$ in the last subsection, its eigenvalues can be efficiently computed for large values of n (say $n \leq 10^6$ on contemporary desktop computers). While computational complexity is no longer such a constraint, accuracy still is: for large n and λ, some eigenvalues of $H_{\lambda,[1,n]}$ are closer than the precision of the floating point arithmetic, rendering, for example, $|E_1 - E_2| = 0$ for theoretically distinct eigenvalues E_1 and E_2 of $H_{\lambda,[1,n]}$.[17]

Figure 9 shows estimates of the IDS based on computations with $n = 10{,}000$ for λ values ranging from the trivial case of no coupling ($\lambda = 0$) to strong coupling ($\lambda = 8$). The fine structure of the spectrum is evident in Figure 10, which repeatedly zooms in upon subsets of the spectrum of the finite section $H_{\lambda,[1,n]}$ for $\lambda = 1$ and $n = 100{,}000$. (The numerical concerns discussed in the last paragraph do not affect these figures.)

We now explore the Hölder continuity of the integrated density of states. In consideration of (14), define

$$N_{n,\lambda}(E) = \lim_{n\to\infty} \frac{\#\{\text{eigenvalues of } H_{\lambda,[1,n]} \text{ that are} \leq E\}}{n}.$$

[17]By its structure, $H_{\lambda,[1,n]}$ must have n distinct eigenvalues. Similar scenarios with exceptionally close distinct eigenvalues are well known in the numerical analysis community; see, e.g., Wilkinson's W_{21}^+ matrix [147, Sec. 7.7].

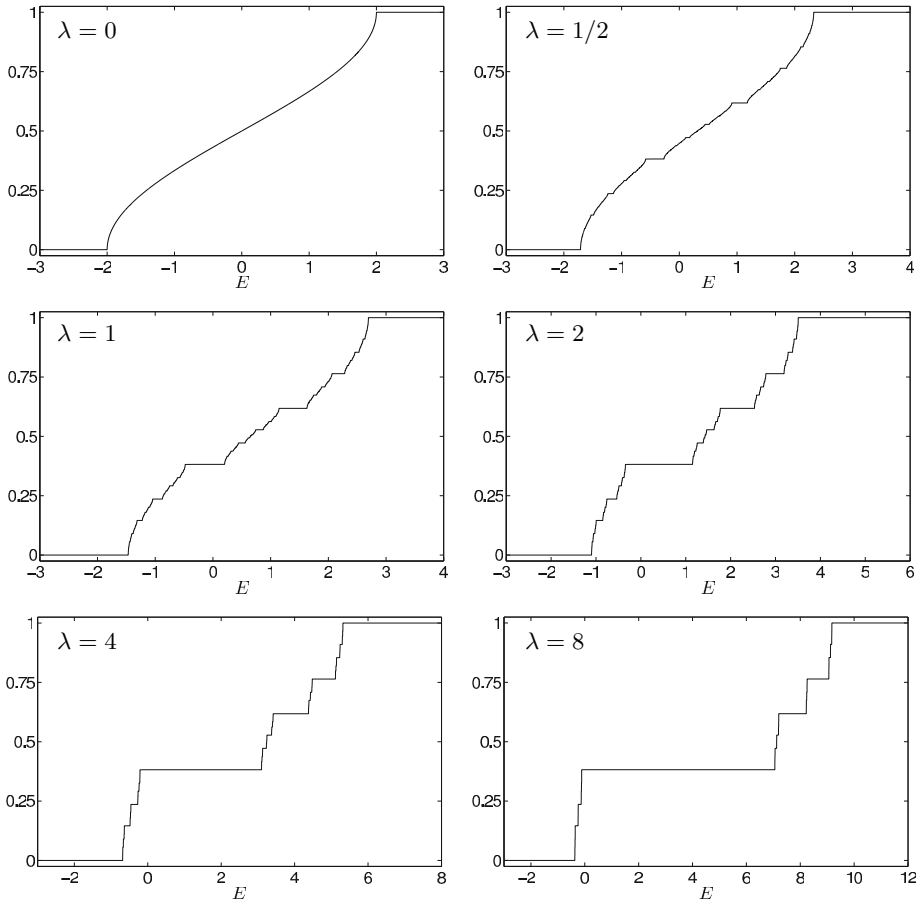

FIGURE 9. Approximations to the integrated density of states for the Fibonacci model with various values of the coupling constant, λ, based on $n = 10{,}000$.

Figure 11 investigates the large λ behavior of the Hölder exponent addressed in Theorem 5.12, based on computations with finite sections of dimension $n = 10{,}000$. Indeed, we see asymptotic behavior like $\frac{3\log(\alpha^{-1})}{2\log\lambda}$, and moreover the figure suggests that the dimension of the measure is smooth in this regime.

7.3. Spectral estimates for square and cubic Fibonacci Hamiltonians

As described in Section 5.9, the estimates $\Sigma_{k,\lambda}$ for the one-dimensional Fibonacci spectrum can readily be translated into approximations for the square and cubic cases, as investigated by Even-Dar Mandel and Lifshitz [79]. As described in Theorem 5.26, Σ_λ need not be a Cantor set, especially for small coupling constants. This behavior is apparent in Figures 12 and 13, which illustrate $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ and

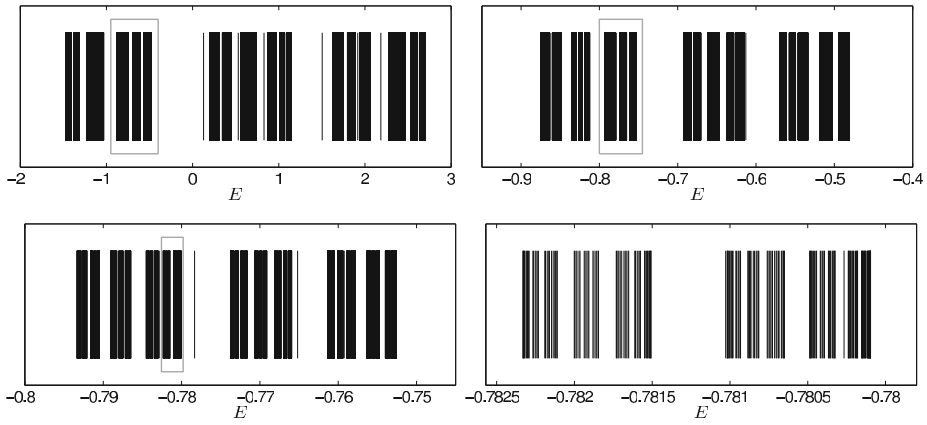

FIGURE 10. Eigenvalues of $H_{\lambda,[1,n]}$ for $\lambda = 1$ and $n = 100{,}000$, drawn as vertical lines to aid visibility. The first plot shows the entire spectrum; the gray boxes denote the region on which the next plot zooms.

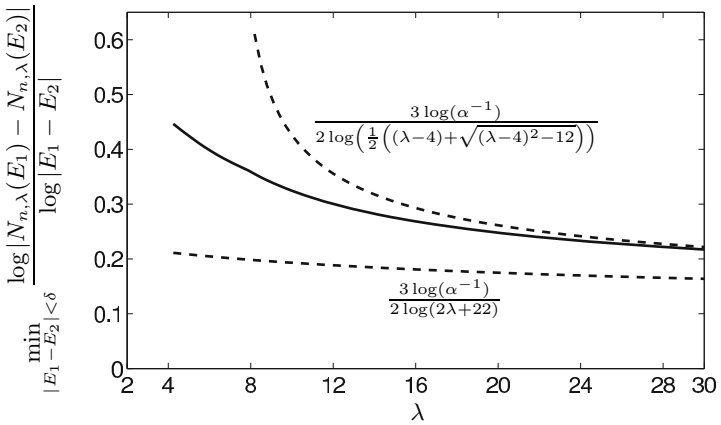

FIGURE 11. Illustration of Theorem 5.12, based on numerically-computed eigenvalues from finite sections $H_{\lambda,[1,n]}$ for $n = 10{,}000$. Here $\delta = 0.025$ and the minimization is over $E_1, E_2 \in \sigma(H_{\lambda,[1,n]})$.

$\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ for various values of k and λ. For a finite range of small λ values, the spectra comprise intervals that branch into a greater number of intervals as k and λ increase. Figure 14 shows the growth in the number of intervals present in these approximations as a function of λ for three different values of k. This plot makes evident rapid (but not always monotone) growth in the number of intervals with λ. Figure 15 illustrates the opening and closing of gaps for the square problem, revealing an intriguing structure for finite k. How does this structure develop

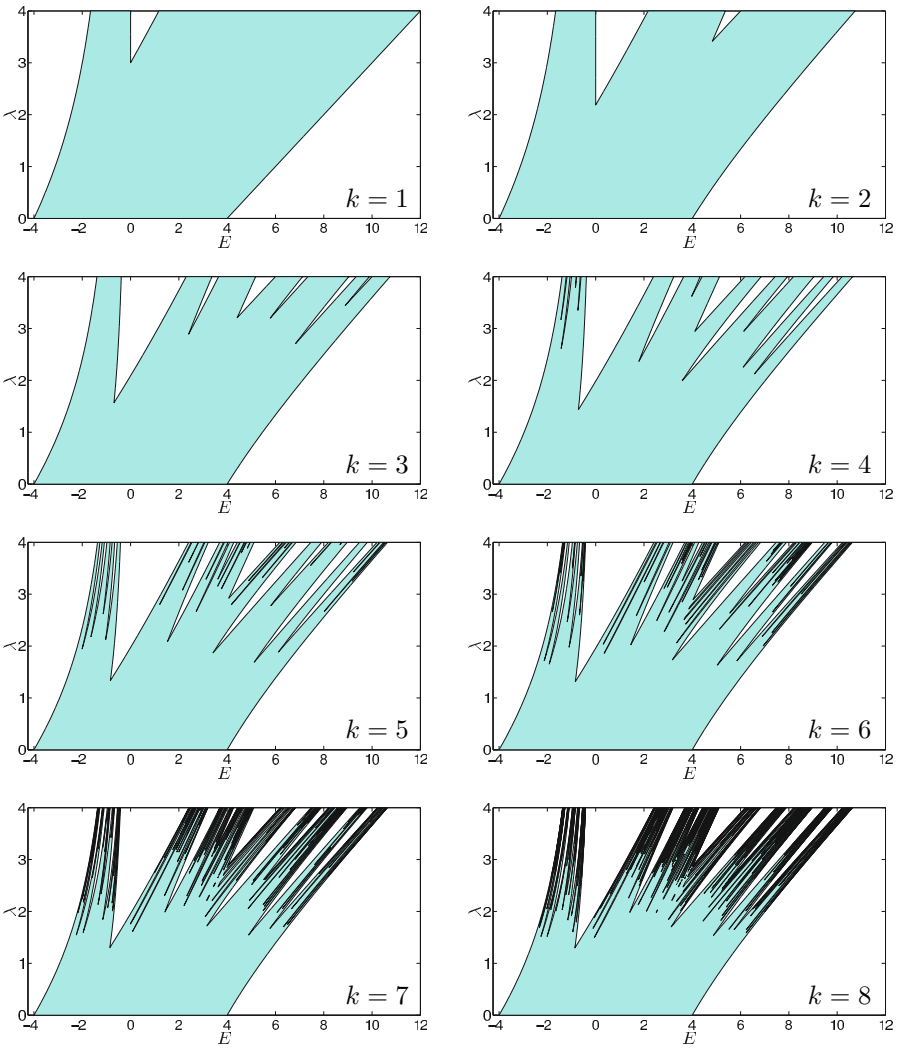

FIGURE 12. Approximations $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ of the spectrum of the square Fibonacci operator, as a function of λ. For $k = 8$ and $\lambda = 4$, $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ is the union of 311 disjoint intervals.

as k increases, and, indeed, is it reflected in $\Sigma_\lambda + \Sigma_\lambda$? At present these questions remain open.

Tables 1 and 2 investigate the square and cubic spectral estimates more precisely, giving the values of λ where multiple intervals first emerge. These results confirm and sharpen the observation of Even-Dar Mandel and Lifshitz [79] that

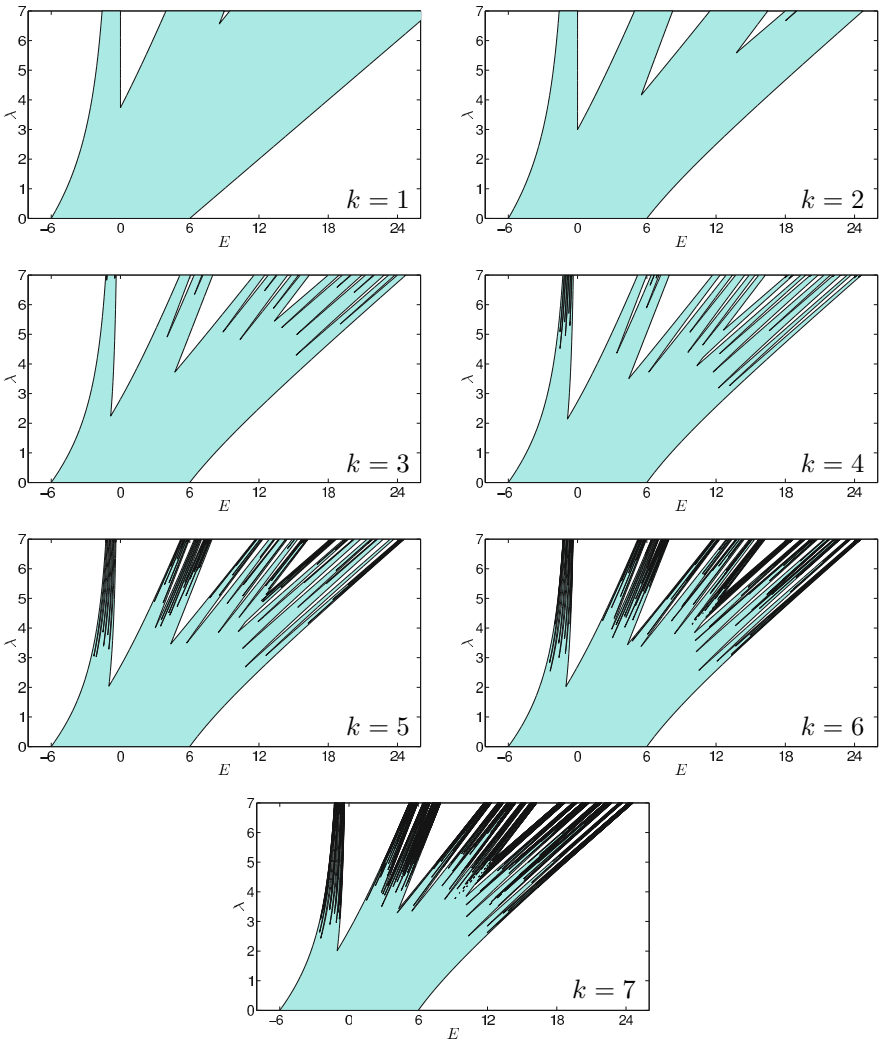

FIGURE 13. Approximations $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ of the spectrum of the cubic Fibonacci operator, as a function of λ. For $k = 7$ and $\lambda = 7$, $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ is the union of 482 disjoint intervals.

$\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ transitions from one to two intervals near $\lambda = 1.3$, while $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ makes the same transition near $\lambda = 2$. For these finite values of k, it is apparent that $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ and $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ both transition to two intervals, then three intervals, and so on. What do these calculations suggest about the limit $k \to \infty$? For example, is the λ value at which $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ transitions from two to

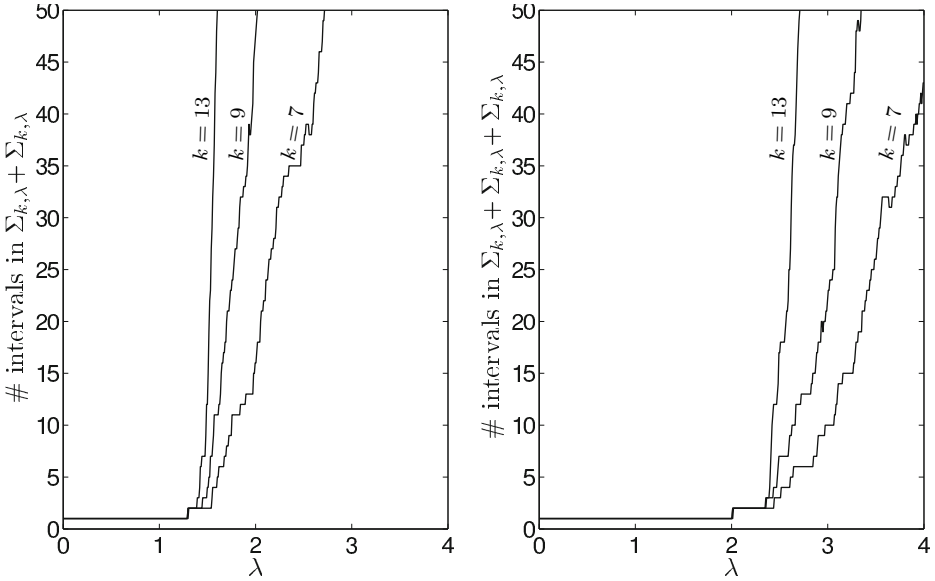

FIGURE 14. Number of intervals in the spectral approximations $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ and $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$, as a function of λ.

three intervals converging? Is there a finite span of λ values for which $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ persists as the union of two intervals as $k \to \infty$, or does $\Sigma_\lambda + \Sigma_\lambda$ transition from one interval directly to a Cantorval or Cantor set? (See Problems 8.6 and 8.7 below.)

Let $\lambda_{k,m}$ denote the value of λ at which $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ (or $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$) first splits from m to $m+1$ intervals as λ increases, with $\lambda_{k,0} = 0$. (Our detailed computations suggest that, for small values of m, there is only one such point of transition; for larger numbers of intervals, gap closings complicate the picture, as seen in Figure 15.) Figure 16 plots $\lambda_{k,m} - \lambda_{k,m-1}$ as a function of k for $m = 1, \ldots, 7$ for the square and cubic Hamiltonians. Do the transition points converge as $k \to \infty$? First consider the plot on the left, for the square Hamiltonian. For $m = 1$ and $m = 2$, $\lambda_{k,m}$ appears to converge; however, the points of transition to $m \geq 3$ intervals do not show such consistency: it is unclear if these $\lambda_{k,m}$ values are converging. It may be that the coupling constants at which $\Sigma_{\lambda,k} + \Sigma_{\lambda,k}$ breaks into $m > 3$ intervals are converging to the point at which the spectrum breaks into $m = 3$ intervals as $k \to \infty$. Now consider the plot on the right of Figure 16, for the cubic Hamiltonian. In contrast to the square case, these results suggest the $\lambda_{k,m}$ values converge to distinct points as $k \to \infty$ for all values $m = 1, 2, \ldots, 7$ shown, inviting the conjecture that there exist λ values for which $\Sigma_\lambda + \Sigma_\lambda + \Sigma_\lambda$ is the union of m disjoint intervals for all $m \geq 1$.

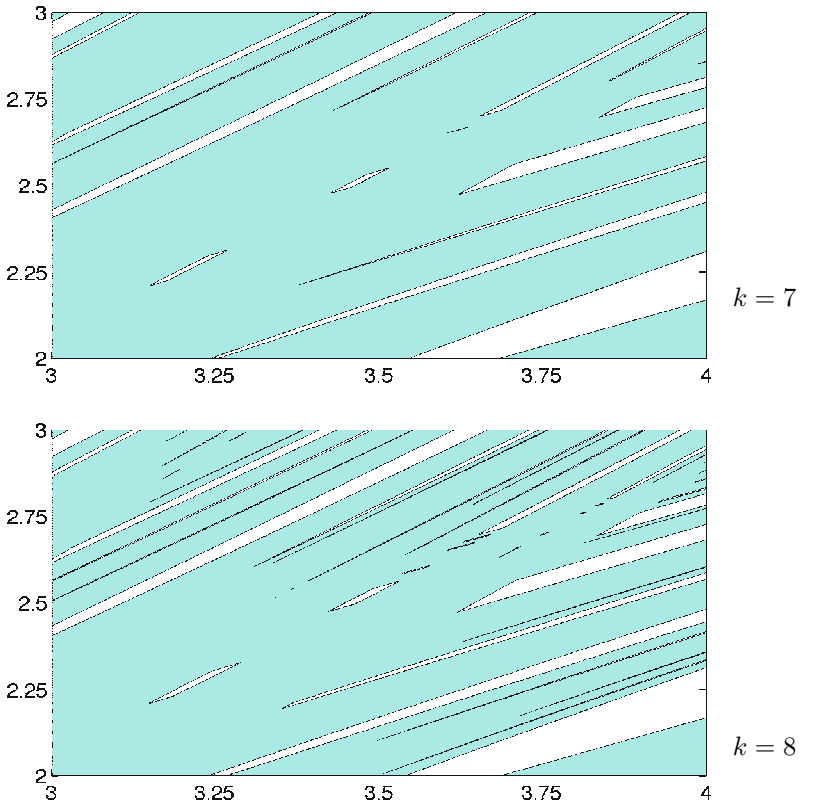

FIGURE 15. Approximations $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ of the square Fibonacci spectrum $\Sigma_\lambda + \Sigma_\lambda$ as in Figure 12, magnified to show the opening and closing of gaps as λ increases. How this structure affects $\Sigma_\lambda + \Sigma_\lambda$ is not presently understood.

8. Conjectures and open problems

In this final section we discuss various open problems that are suggested by the existing results and address generalizations, strengthenings, and related issues.

We begin with open problems for the Fibonacci Hamiltonian. The existing quantitative results concern estimates for dimensional properties of the spectrum, the density of states measure, and the spectral measures, as well as estimates for the transport exponents. In almost all cases, the asymptotic behavior is known in the regimes of small and large coupling. While the bounds we obtain are monotone, we would like to understand whether the quantities themselves have this property:

Problem 8.1. Are the various quantities we consider (in particular, $\dim_H \Sigma_\lambda$) monotone in λ?

TABLE 1. Estimates of λ^*, the λ value for which the thickness of $\Sigma_{k,\lambda}$ equals one, along with $\lambda_{k,m}$, the coupling constant where $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ splits from m to $m+1$ intervals, for $m = 1, \ldots, 4$.

k	λ_k^*	$\lambda_{k,1}$	$\lambda_{k,2}$	$\lambda_{k,3}$	$\lambda_{k,4}$
6	1.313172936	1.313172936	1.624865906	1.649775155	1.708521471
7	1.298964798	1.298964798	1.543759898	1.548912772	1.596682038
8	1.296218739	1.296218739	1.494856217	1.514291562	1.520122025
9	1.294303086	1.294303086	1.445808095	1.492410878	1.512965310
10	1.293935333	1.293935333	1.442778219	1.446787662	1.472813609
11	1.293679331	1.293679331	1.430901095	1.436192692	1.437915282
12	1.293630242	1.293630242	1.402035016	1.415460742	1.426586813
13	1.290031553	1.293596081	1.392730451	1.412863780	1.419815054
14	1.288819456	1.293589532	1.382510414	1.404399139	1.408704405
15	1.287431935	1.293584975	1.380466052	1.399646887	1.400190389
16	1.287269802	1.293584102	1.380121550	1.388518687	1.397593470
17	1.287084388	1.293583494	1.379851608	1.387310733	1.395556145
18	1.287062735	1.293583377	1.379806139	1.385835331	1.393702258
19	1.287037977				
20	1.287035086				

TABLE 2. Estimates of $\lambda_{k,m}$, the coupling constant where $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ splits from m to $m+1$ intervals, for $m = 1, \ldots, 4$.

k	$\lambda_{k,1}$	$\lambda_{k,2}$	$\lambda_{k,3}$	$\lambda_{k,4}$
6	2.025741216	2.544063632	2.573539294	2.842670115
7	2.012664501	2.438240772	2.511570744	2.606841186
8	2.011113604	2.376933028	2.498126298	2.498926850
9	2.009524869	2.364541039	2.435665993	2.473875055
10	2.009337409	2.357357667	2.412613336	2.421115367
11	2.009145619	2.355932060	2.399696274	2.408616763
12	2.009123008	2.355107791	2.392573154	2.401253561
13	2.009099880	2.354944739	2.391094663	2.397036745
14	2.009097154	2.354850520	2.390282080	2.393347062
15	2.009094365	2.354831891	2.390113912	2.393329303
16	2.009094036	2.354821128	2.390021550	2.393302392
17	2.009093700	2.354819000	2.390002443	2.393300376

The known estimates for the local scaling exponents and in particular the optimal Hölder exponent of the spectral measures (see [52] and references therein) are clearly not optimal, and in particular do not identify their asymptotics in the extremal coupling regimes. For the density of states measure, which is an average

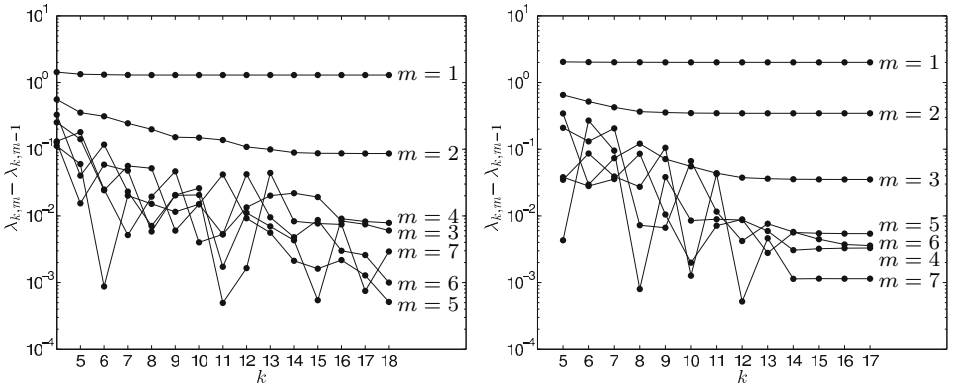

FIGURE 16. The span of λ values (i.e., $\lambda_{k,m} - \lambda_{k,m-1}$) for which $\Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ (left) and $\Sigma_{k,\lambda} + \Sigma_{k,\lambda} + \Sigma_{k,\lambda}$ (right) comprise m intervals for $m = 1, \ldots, 7$.

of spectral measures, we have much better information. Can one find ways to find equally good estimates for spectral measures?

Problem 8.2. What can one say about the spectral measures? In particular, are their dimensional properties uniform across the hull and/or across the spectrum? Moreover, what are the asymptotics as $\lambda \downarrow 0$ and $\lambda \uparrow \infty$?

We know that $\dim \Sigma_\lambda$ goes to one as λ goes to zero. In addition, we would be interested in the following:

Problem 8.3. Does the right-derivative of $\dim \Sigma_\lambda$ exist at zero?

If it does, due to Theorem 5.9 it must be finite and non-zero.

Let us now turn to the higher-dimensional separable analogs of the Fibonacci Hamiltonian (e.g., the square or cubic Fibonacci Hamiltonian). Recall that the spectrum of such an operator is given by the sum of the one-dimensional spectra, which in turn are Cantor sets. Recall also that at sufficiently small coupling, these sum sets are intervals, while at sufficiently large coupling, they are Cantor sets as well. Concretely, this uses that if the thickness of a Cantor set C is larger than 1, then $C + C$ is an interval by Theorem 5.4 and, on the other hand, if the upper box counting dimension of C is strictly less than $1/2$, then $C + C$ is a Cantor set. It is natural to ask what shape the higher-dimensional spectra have at intermediate coupling, that is, we wish to study how the transition from $C + C$ being an interval to being a Cantor set happens when the thickness of C decreases.

Definition 8.4. A compact set $C \subset \mathbb{R}^1$ is a Cantorval if it has a dense interior (i.e., $\overline{\text{int}(C)} = C$), it has a continuum of connected components, and none of them is isolated.

Here is a general result on the occurrence of Cantorvals in the context of taking sums of Cantor sets [138]:

Theorem 8.5. *There is an open set $\mathcal{U}$ in the space of dynamically defined Cantor sets such that for generic $C_1, C_2 \in \mathcal{U}$, the sum $C_1 + C_2$ is a Cantorval.*

Unfortunately, this result does not provide any specific and verifiable genericity conditions that would allow one to check that the sum of two given specific Cantor sets is indeed a Cantorval. Thus, for our purpose we need a solution to the following problem.

Problem 8.6. Provide specific verifiable conditions on a Cantor set C which imply that the sum $C + C$ is a Cantorval.

Ideally, such a criterion would be applicable to the spectrum of the Fibonacci Hamiltonian and establish that, say, the spectrum of the square Fibonacci Hamiltonian is a Cantorval for intermediate values of the coupling constant λ. The next step would then be to study the transitions between the three regimes. We ask whether there are two sharp transitions; compare [80] for closely related numerical evidence and discussion.

Problem 8.7. Let $H_\lambda^{(2)}$ be the separable square Fibonacci Hamiltonian. Prove that there are values $0 < \lambda' < \lambda'' < \infty$ such that for $\lambda \in (0, \lambda')$, the spectrum $\sigma(H_\lambda)$ is an interval (or a finite union of intervals), for $\lambda \in (\lambda', \lambda'')$, it is a Cantorval, and for $\lambda \in (\lambda'', \infty)$, it is a Cantor set.

Notice that this will provide an example of a (*topologically!*) new structure of the spectrum for "natural" potentials.

Moving on from the Fibonacci case, which has a description via a substitution rule as well as via a simple quasi-periodic expression, there are two natural choices of a more general setting.

For a different choice of the underlying substitution rule, one always has an associated trace map. However, our understanding of the dynamics of such a trace map is in general far more limited than the one in the Fibonacci case. As a consequence, outside of the Fibonacci case there is a scarcity of quantitative results for dimensional issues (such as the dimension of the spectrum, the dimension of the density of states measure, or the dimension of the spectral measures). For example, here is a simple open problem that is currently completely out of reach:

Problem 8.8. Study other trace maps (e.g., period doubling and Thue–Morse); in particular, find the asymptotics of the Hausdorff dimension of the spectrum as the coupling constant tends to zero and infinity.

The other natural generalization of the Fibonacci potential is to replace the golden ratio in its quasi-periodic description by a general irrational number. Thus, we discuss some open problems for Sturmian potentials next.

Let us say that two Cantor sets C_1 and C_2 on $\mathbb{R}^1$ are *diffeomorphic* if there are neighborhoods $U_1(C_1)$, $U_2(C_2)$, and a diffeomorphism $f : U_1 \to U_2$ such that $f(C_1) = C_2$.

Problem 8.9. Suppose that $\alpha = [a_1, a_2, \ldots]$ and $\beta = [b_1, b_2, \ldots]$ are such that for some $k \in \mathbb{Z}$ and all large enough $i \in \mathbb{Z}_+$ we have $b_{i+k} = a_i$. Prove that in this case, the Sturmian spectra $\Sigma_{\lambda,\alpha}$ and $\Sigma_{\lambda,\beta}$ are diffeomorphic.

Notice also that due to the ergodicity of the Gauss map, a solution of this problem would also imply that the following long standing conjecture is correct:

Problem 8.10. For any fixed $\lambda > 0$, the dimension $\dim_H \Sigma_{\lambda,\alpha}$ is almost everywhere constant in α.

Finally, let us emphasize that most of the questions related to higher-dimensional models (described in Section 3) are completely open. So we formulate an extremely general problem:

Problem 8.11. Study spectral properties (e.g., the shape of the spectrum and the type of the spectral measures) and transport properties of higher-dimensional operators; for example, study these questions for the particular case of the Laplacian on the graph associated with a Penrose tiling.

References

[1] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, *LAPACK User's Guide*, 3rd ed., SIAM, Philadelphia, 1999.

[2] J. Ashraff, R. Stinchcombe, Dynamic structure factor for the Fibonacci-chain quasicrystal, *Phys. Rev. B* **39** (1989), 2670–2677.

[3] Y. Avishai, D. Berend, Trace maps for arbitrary substitution sequences, *J. Phys. A* **26** (1993), 2437–2443.

[4] A. Avila, S. Jitomirskaya, The Ten Martini Problem, *Ann. of Math.* **170** (2009), 303–342.

[5] A. Avila, S. Jitomirskaya, Almost localization and almost reducibility, *J. Eur. Math. Soc.* (*JEMS*) **12** (2010), 93–131.

[6] M. Baake, U. Grimm, The singular continuous diffraction measure of the Thue–Morse chain, *J. Phys. A* **41** (2008), no. 42, 422001.

[7] M. Baake, U. Grimm, D. Joseph, Trace maps, invariants, and some of their applications, *Internat. J. Modern Phys. B* **7** (1993), 1527–1550.

[8] L. Barreira, Ya. Pesin, J. Schmeling, Dimension and product structure of hyperbolic measures, *Ann. of Math.* **149** (1999), 755–783.

[9] L. Barreira, *Thermodynamic Formalism and Applications to Dimension Theory*, Progress in Mathematics **294**, Birkhäuser/Springer Basel AG, Basel, 2011.

[10] J. Bellissard, Spectral properties of Schrödinger's operator with a Thue–Morse potential, *Number Theory and Physics* (*Les Houches,* 1989), 140–150, Springer Proc. Phys. **47**, Springer, Berlin, 1990.

[11] J. Bellissard, Renormalization group analysis and quasicrystals, *Ideas and Methods in Quantum and Statistical Physics* (*Oslo,* 1988), 118–148, Cambridge Univ. Press, Cambridge, 1992.

[12] J. Bellissard, A. Bovier, J.-M. Ghez, Spectral properties of a tight binding Hamiltonian with period doubling potential, *Commun. Math. Phys.* **135** (1991), 379–399.

[13] J. Bellissard, A. Bovier, J.-M. Ghez, Gap labelling theorems for one-dimensional discrete Schrödinger operators, *Rev. Math. Phys.* **4** (1992), 1–37.

[14] J. Bellissard, B. Iochum, E. Scoppola, D. Testard, Spectral properties of one-dimensional quasicrystals, *Commun. Math. Phys.* **125** (1989), 527–543.

[15] V. Benza, Quantum Ising quasi-crystal, *Europhys. Lett.* **8** (1989), 321–325.

[16] J. Breuer, Y. Last, Y. Strauss, Eigenvalue spacings and dynamical upper bounds for discrete one-dimensional Schrödinger operators, *Duke Math. J.* **157** (2011), 425–460.

[17] R. Brown, The algebraic entropy of the special linear character automorphisms of a free group on two generators, *Trans. Amer. Math. Soc.* **359** (2007), 1445–1470.

[18] J. Berstel, Recent results in Sturmian words, in *Developments in Language Theory*, eds. J. Dassow and A. Salomaa, World Scientific, Singapore (1996), 13–24.

[19] B. Bowditch, A proof of McShane's identity via Markoff triples, *Bull. London Math. Soc.* **28** (1996), 73–78.

[20] B. Bowditch, Markoff triples and quasi-Fuchsian groups, *Proc. London Math. Soc.* **77** (1998), 697–736.

[21] M. Boshernitzan, A condition for minimal interval exchange maps to be uniquely ergodic, *Duke Math. J.* **52** (1985), 723–752.

[22] M. Boshernitzan, Rank two interval exchange transformations, *Ergod. Th. & Dynam. Sys.* **8** (1988), 379–394.

[23] J. Bourgain, Hölder regularity of integrated density of states for the almost Mathieu operator in a perturbative regime, *Lett. Math. Phys.* **51** (2000), 83–118.

[24] J. Bourgain, Green's function estimates for lattice Schrödinger operators and applications, *Ann. Math. Stud.* **158**, Princeton University Press, Princeton, NJ (2005).

[25] J. Bourgain, M. Goldstein, W. Schlag, Anderson localization for Schrödinger operators on $\mathbb{Z}$ with potentials given by the skew-shift, *Commun. Math. Phys.* **220** (2001), 583–621.

[26] A. Bovier, J.-M. Ghez, Spectral properties of one-dimensional Schrödinger operators with potentials generated by substitutions, *Commun. Math. Phys.* **158** (1993), 45–66; Erratum: *Commun. Math. Phys.* **166** (1994), 431–432.

[27] A. Bovier, J.-M. Ghez, Remarks on the spectral properties of tight-binding and Kronig–Penney models with substitution sequences, *J. Phys. A* **28** (1995), 2313–2324.

[28] R. Bowen, *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*, Lect. Notes in Math. **470**, Springer (1975).

[29] M. Campanino, A. Klein, A supersymmetric transfer matrix and differentiability of the density of states in the one-dimensional Anderson model, *Commun. Math. Phys.* **104** (1986), 227–241.

[30] S. Cantat, Bers and Hénon, Painlevé and Schrödinger, *Duke Math. J.* **149** (2009), 411–460.

[31] S. Cantat, F. Loray, Dynamics on character varieties and Malgrange irreducibility of Painlevé VI equation, *Ann. Inst. Fourier* (*Grenoble*) **59** (2009), 2927–2978.

[32] R. Carmona, J. Lacroix, *Spectral Theory of Random Schrödinger Operators*, Probability and its Applications, Birkhäuser Boston, Inc., Boston, MA, 1990.

[33] M. Casdagli, Symbolic dynamics for the renormalization map of a quasiperiodic Schrödinger equation, *Comm. Math. Phys.* **107** (1986), 295–318.

[34] J. Cassaigne, Sequences with grouped factors, in *Developments in Language Theory III*, Aristotle University of Thessaloniki (1998), 211–222.

[35] M. Cobo, C. Gutierrez, C. de Oliveira, Cantor singular continuous spectrum for operators along interval exchange transformations, *Proc. Amer. Math. Soc.* **136** (2008), 923–930.

[36] E. Coven, Sequences with minimal block growth, II, *Math. Systems Theory* **8** (1975), 376–382.

[37] W. Craig, Pure point spectrum for discrete almost periodic Schrödinger operators, *Commun. Math. Phys.* **88** (1983), 113–131.

[38] W. Craig, B. Simon, Subharmonicity of the Lyaponov index, *Duke Math. J.* **50** (1983), 551–560.

[39] W. Craig, B. Simon, Log Hölder continuity of the integrated density of states for stochastic Jacobi matrices, *Commun. Math. Phys.* **90** (1983), 207–218.

[40] D. Damanik, α-continuity properties of one-dimensional quasicrystals, *Commun. Math. Phys.* **192** (1998), 169–182.

[41] D. Damanik, Singular continuous spectrum for the period doubling Hamiltonian on a set of full measure, *Commun. Math. Phys.* **196** (1998), 477–483.

[42] D. Damanik, Singular continuous spectrum for a class of substitution Hamiltonians. II, *Lett. Math. Phys.* **54** (2000), 25–31.

[43] D. Damanik, Gordon-type arguments in the spectral theory of one-dimensional quasicrystals, in *Directions in Mathematical Quasicrystals*, CRM Monogr. Ser. **13**, Amer. Math. Soc., Providence, RI (2000), 277–305.

[44] D. Damanik, Uniform singular continuous spectrum for the period doubling Hamiltonian, *Ann. Henri Poincaré* **2** (2001), 101–108.

[45] D. Damanik, Dynamical upper bounds for one-dimensional quasicrystals, *J. Math. Anal. Appl.* **303** (2005), 327–341.

[46] D. Damanik, Strictly ergodic subshifts and associated operators, in *Spectral Theory and Mathematical Physics: a Festschrift in Honor of Barry Simon's* 60*th Birthday*, Proc. Sympos. Pure Math. **76**, Part 2, Amer. Math. Soc., Providence, RI (2007), 505–538.

[47] D. Damanik, Lyapunov exponents and spectral analysis of ergodic Schrödinger operators: a survey of Kotani theory and its applications, in *Spectral Theory and Mathematical Physics: a Festschrift in Honor of Barry Simon's* 60*th Birthday*, Proc. Sympos. Pure Math. **76**, Part 2, Amer. Math. Soc., Providence, RI (2007), 539–563.

[48] D. Damanik, M. Embree, A. Gorodetski, S. Tcheremchantsev, The fractal dimension of the spectrum of the Fibonacci Hamiltonian, *Commun. Math. Phys.* **280** (2008), 499–516.

[49] D. Damanik, J. Fillman, M. Lukic, W. Yessen, Uniform hyperbolicity for Szegő cocycles and applications to random CMV matrices and the Ising model, preprint (arXiv:1404.7065).

[50] D. Damanik, A. Gorodetski, Hyperbolicity of the trace map for the weakly coupled Fibonacci Hamiltonian, *Nonlinearity* **22** (2009), 123–143.

[51] D. Damanik, A. Gorodetski, The spectrum of the weakly coupled Fibonacci Hamiltonian, *Electronic Research Announcements in Mathematical Sciences*, **16** (2009), 23–29.

[52] D. Damanik, A. Gorodetski, Spectral and quantum dynamical properties of the weakly coupled Fibonacci Hamiltonian, *Commun. Math. Phys.* **305** (2011), 221–277.

[53] D. Damanik, A. Gorodetski, The density of states measure of the weakly coupled Fibonacci Hamiltonian, *Geom. Funct. Anal.* **22** (2012), 976–989.

[54] D. Damanik, A. Gorodetski, Hölder continuity of the integrated density of states for the Fibonacci Hamiltonian, *Commun. Math. Phys.* **323** (2013), 497–515.

[55] D. Damanik, A. Gorodetski, Almost ballistic transport for the weakly coupled Fibonacci Hamiltonian, to appear in *Israel J. Math.* (arXiv:1307.0925).

[56] D. Damanik, A. Gorodetski, B. Solomyak, Absolutely continuous convolutions of singular measures and an application to the Square Fibonacci Hamiltonian, preprint (arXiv:1306.4284).

[57] D. Damanik, A. Gorodetski, W. Yessen, The Fibonacci Hamiltonian, preprint (arXiv:1403.7823).

[58] D. Damanik, R. Killip, D. Lenz, Uniform spectral properties of one-dimensional quasicrystals. III. α-continuity, *Commun. Math. Phys.* **212** (2000), 191–204.

[59] D. Damanik, D. Lenz, Uniform spectral properties of one-dimensional quasicrystals, I. Absence of eigenvalues, *Commun. Math. Phys.* **207** (1999), 687–696.

[60] D. Damanik, D. Lenz, Uniform spectral properties of one-dimensional quasicrystals. II. The Lyapunov exponent, *Lett. Math. Phys.* **50** (1999), 245–257.

[61] D. Damanik, D. Lenz, Uniform spectral properties of one-dimensional quasicrystals, IV. Quasi-Sturmian potentials, *J. Anal. Math.* **90** (2003), 115–139.

[62] D. Damanik, D. Lenz, A condition of Boshernitzan and uniform convergence in the multiplicative ergodic theorem, *Duke Math. J.* **133** (2006), 95–123.

[63] D. Damanik, D. Lenz, Zero-measure Cantor spectrum for Schrödinger operators with low-complexity potentials, *J. Math. Pures Appl.* (9) **85** (2006), 671–686.

[64] D. Damanik, D. Lenz, Substitution dynamical systems: Characterization of linear repetitivity and applications, *J. Math. Anal. Appl.* **321** (2006), 766–780.

[65] D. Damanik, P. Munger, W. Yessen, Orthogonal polynomials on the unit circle with Fibonacci Verblunsky coefficients, I. The essential support of the measure, *J. Approx. Theory* **173** (2013), 56–88.

[66] D. Damanik, A. Sütő, S. Tcheremchantsev, Power-law bounds on transfer matrices and quantum dynamics in one dimension II., *J. Funct. Anal.* **216** (2004), 362–387.

[67] D. Damanik, S. Tcheremchantsev, Power-law bounds on transfer matrices and quantum dynamics in one dimension, *Commun. Math. Phys.* **236** (2003), 513–534.

[68] D. Damanik, S. Tcheremchantsev, Scaling estimates for solutions and dynamical lower bounds on wavepacket spreading, *J. d'Analyse Math.* **97** (2005), 103–131.

[69] D. Damanik, S. Tcheremchantsev, Upper bounds in quantum dynamics, *J. Amer. Math. Soc.* **20** (2007), 799–827.

[70] D. Damanik, S. Tcheremchantsev, Quantum dynamics via complex analysis methods: general upper bounds without time-averaging and tight lower bounds for the strongly coupled Fibonacci Hamiltonian, *J. Funct. Anal.* **255** (2008), 2872–2887.

[71] D. Damanik, S. Tcheremchantsev, A general description of quantum dynamical spreading over an orthonormal basis and applications to Schrödinger operators, *Discrete Contin. Dyn. Syst.* **28** (2010), 1381–1412.

[72] D. Damanik, D. Zare, Palindrome complexity bounds for primitive substitution sequences, *Discrete Math.* **222** (2000), 259–267.

[73] F. Delyon, D. Petritis, Absence of localization in a class of Schrödinger operators with quasiperiodic potential, *Commun. Math. Phys.* **103** (1986), 441–444.

[74] C. de Oliveira, C. Gutierrez, Almost periodic Schrödinger operators along interval exchange transformations, *J. Math. Anal. Appl.* **283** (2003), 570–581.

[75] C. de Oliveira, M. Lima, A nonprimitive substitution Schrödinger operator with generic singular continuous spectrum, *Rep. Math. Phys.* **45** (2000), 431–436.

[76] C. de Oliveira, M. Lima, Singular continuous spectrum for a class of nonprimitive substitution Schrödinger operators, *Proc. Amer. Math. Soc.* **130** (2002), 145–156.

[77] P. Duarte, Persistent homoclinic tangencies for conservative maps near the identity, *Ergod. Th. & Dynam. Sys.* **20** (2000), 393–438.

[78] F. Durand, B. Host, C. Skau, Substitution dynamical systems, Bratteli diagrams and dimension groups, *Ergod. Th. & Dynam. Sys.* **19** (1999), 953–993.

[79] S. Even-Dar Mandel, R. Lifshitz, Electronic energy spectra and wave functions on the square Fibonacci tiling, *Phil. Mag.* **86** (2006), 759–764.

[80] S. Even-Dar Mandel, R. Lifshitz, Electronic energy spectra of square and cubic Fibonacci quasicrystals, *Phil. Mag.* **88** (2008), 2261–2273.

[81] S. Even-Dar Mandel, R. Lifshitz, Bloch-like electronic wave functions in two-dimensional quasicrystals, Preprint (arXiv:0808.3659).

[82] S. Fan, Q.-H. Liu, Z.-Y. Wen, Gibbs-like measure for spectrum of a class of quasicrystals, *Ergod. Th. & Dynam. Sys.* **31** (2011), 1669–1695.

[83] Z. Gan, H. Krüger, Optimality of log Hölder continuity of the integrated density of states, *Math. Nachr.* **284** (2011), 1919–1923.

[84] G. Gumbs, M. Ali, Electronic properties of the tight-binding Fibonacci Hamiltonian, *J. Phys. A* **22** (1989), 951–970.

[85] W. Goldman, Mapping class group dynamics on surface group representations, in *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, 189–214, Amer. Math. Soc., Providence, RI, 2006.

[86] M. Goldstein, W. Schlag, Hölder continuity of the integrated density of states for quasi-periodic Schrödinger equations and averages of shifts of subharmonic functions, *Ann. of Math.* **154** (2001), 155–203.

[87] M. Goldstein, W. Schlag, Fine properties of the integrated density of states and a quantitative separation property of the Dirichlet eigenvalues, *Geom. Funct. Anal.* **18** (2008), 755–869.

[88] A. Gordon, On the point spectrum of the one-dimensional Schrödinger operator, *Usp. Math. Nauk.* **31** (1976), 257–258.

[89] S. Hadj Amor, Hölder continuity of the rotation number for quasi-periodic co-cycles in SL(2, $\mathbb{R}$), *Commun. Math. Phys.* **287** (2009), 565–588.

[90] G. Hedlund, M. Morse, Symbolic dynamics, *Amer. J. Math.* **60** (1938), 815–866.

[91] J. Hermisson, U. Grimm, M. Baake, Aperiodic Ising quantum chains, *J. Phys. A* **30** (1997), 7315–7335.

[92] M. Hirsch, C. Pugh, M. Shub, *Invariant Manifolds*, Lecture Notes on Mathematics **583**, Springer, Heidelberg, 1977.

[93] A. Hof, Some remarks on discrete aperiodic Schrödinger operators, *J. Statist. Phys.* **72** (1993), 1353–1374.

[94] A. Hof, A remark on Schrödinger operators on aperiodic tilings, *J. Statist. Phys.* **81** (1995), 851–855.

[95] A. Hof, O. Knill, and B. Simon, Singular continuous spectrum for palindromic Schrödinger operators, *Commun. Math. Phys.* **174** (1995), 149–159.

[96] M. Holzer, Three classes of one-dimensional, two-tile Penrose tilings and the Fibonacci Kronig–Penney model as a generic case, *Phys. Rev. B* **38** (1988), 1709–1720.

[97] M. Holzer, Nonlinear dynamics of localization in a class of one-dimensional quasicrystals, *Phys. Rev. B* **38** (1988), 5756–5759.

[98] S. Humphries, A. Manning, Curves of fixed points of trace maps, *Ergod. Th. & Dynam. Sys.*, **27** (2007), 1167–1198.

[99] B. Iochum, L. Raymond, D. Testard, Resistance of one-dimensional quasicrystals, *Phys. A* **187** (1992), 353–368.

[100] B. Iochum, D. Testard, Power law growth for the resistance in the Fibonacci model, *J. Stat. Phys.* **65** (1991), 715–723.

[101] G. Iommi, The Bowen formula: Dimension theory and thermodynamic formalism, preprint.

[102] K. Ishii, Localization of eigenstates and transport phenomena in the one dimensional disordered system, *Supp. Theor. Phys.* **53** (1973), 77–138.

[103] K. Iwasaki, T. Uehara, An ergodic study of Painlevé VI, *Math. Ann.* **338** (2007), 295–345.

[104] S. Jitomirskaya and Y. Last, Power-law subordinacy and singular spectra. II. Line operators, *Commun. Math. Phys.* **211** (2000), 643–658.

[105] R. Johnson, Exponential dichotomy, rotation number, and linear differential operators with bounded coefficients, *J. Differential Equations* **61** (1986), 54–78.

[106] M. Kaminaga, Absence of point spectrum for a class of discrete Schrödinger operators with quasiperiodic potential, *Forum Math.* **8** (1996), 63–69.

[107] A. Katok, B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Encyclopedia of Mathematics and its Applications **54**, Cambridge University Press, Cambridge, 1995.

[108] M. Keane, Interval exchange transformations, *Math. Z.* **141** (1975), 25–31.

[109] G. Keller, *Equilibrium States in Ergodic Theory*, London Mathematical Society Student Texts **42**, Cambridge University Press, Cambridge, 1998.

[110] R. Killip, A. Kiselev, Y. Last, Dynamical upper bounds on wavepacket spreading, *Amer. J. Math.* **125** (2003), 1165–1198.

[111] S. Klassert, D. Lenz, P. Stollmann, Discontinuities of the integrated density of states for random operators on Delone sets, *Commun. Math. Phys.* **241** (2003), 235–243.

[112] M. Kohmoto, L.P. Kadanoff, C. Tang, Localization problem in one dimension: mapping and escape, *Phys. Rev. Lett.* **50** (1983), 1870–1872.

[113] M. Kohmoto, Dynamical system related to quasiperiodic Schrödinger equations in one dimension, *J. Statist. Phys.* **66** (1992), 791–796.

[114] M. Kohmoto, B. Sutherland, C. Tang, Critical wave functions and a Cantor-set spectrum of a one-dimensional quasicrystal model, *Phys. Rev. B* **35** (1987), 1020–1033.

[115] S. Kotani, Ljapunov indices determine absolutely continuous spectra of stationary random one-dimensional Schrödinger operators, in *Stochastic Analysis* (*Katata/Kyoto,* 1982), North Holland, Amsterdam (1984), 225–247.

[116] S. Kotani, Jacobi matrices with random potentials taking finitely many values, *Rev. Math. Phys.* **1** (1989), 129–133.

[117] M.P. Lamoureux, Reflections on the almost Mathieu operator, *Int. Eqns. Operator Theory* **28** (1997), 45–59.

[118] Y. Last, Quantum dynamics and decompositions of singular continuous spectra, *J. Funct. Anal.* **142** (1996), 406–445.

[119] F. Ledrappier, L.-S. Young, The metric entropy of diffeomorphisms, I. Characterization of measures satisfying Pesin's entropy formula, *Ann. of Math.* **122** (1985), 509–539; The metric entropy of diffeomorphisms, II. Relations between entropy, exponents and dimension, *Ann. of Math.* **122** (1985), 540–574.

[120] D. Lenz, Uniform ergodic theorems on subshifts over a finite alphabet, *Ergod. Th. & Dynam. Sys.* **22** (2002), 245–255.

[121] D. Lenz, Singular continuous spectrum of Lebesgue measure zero for one-dimensional quasicrystals, *Commun. Math. Phys.* **227** (2002), 119–130.

[122] D. Lenz, P. Stollmann, Delone dynamical systems and associated random operators, in *Operator Algebras and Mathematical Physics*, 267–285, Theta, Bucharest, 2003.

[123] D. Lenz, P. Stollmann, An ergodic theorem for Delone dynamical systems and existence of the integrated density of states, *J. Anal. Math.* **97** (2005), 1–24.

[124] D. Levine, Quasicrystals, *J. Physique* **46** (1985), no. 12, Suppl. Colloq. C8, 397–402.

[125] E. Le Page, State distribution of a random Schrödinger operator. Empirical distribution of the eigenvalues of a Jacobi matrix, in *Probability Measures on Groups, VII* (Oberwolfach, 1983), Lecture Notes in Math. **1064**, Springer-Verlag, Berlin (1984), 309–367.

[126] M. Lima, C. de Oliveira, Uniform Cantor singular continuous spectrum for non-primitive Schrödinger operators, *J. Statist. Phys.* **112** (2003), 357–374.

[127] Q.-H. Liu, Y.-H. Qu, Uniform convergence of Schrödinger cocycles over simple Toeplitz subshift, *Ann. Henri Poincaré* **12** (2011), 153–172.

[128] Q.-H. Liu, Y.-H. Qu, Z.-Y. Wen, The fractal dimensions of the spectrum of Sturm Hamiltonian, *Adv. Math.* **257** (2014), 285–336.

[129] Q.-H. Liu, B. Tan, Z.-X. Wen, J. Wu, Measure zero spectrum of a class of Schrödinger operators, *J. Statist. Phys.* **106** (2002), 681–691.

[130] Q.-W. Liu, Z.-Y. Wen, Hausdorff dimension of spectrum of one-dimensional Schrödinger operator with Sturmian potentials, *Potential Anal.* **20** (2004), 33–59.

[131] J. Luck, Cantor spectra and scaling of gap widths in deterministic aperiodic systems, *Phys. Rev. B* **39** (1989), 5834–5849.

[132] N. Makarov, Fine structure of harmonic measure, *St. Petersburg Math. J.* **10** (1999), 217–268.

[133] N. Makarov, A. Volberg, On the harmonic measure of discontinuous fractals, preprint (1986).

[134] R. Mane, The Hausdorff dimension of horseshoes of diffeomorphisms of surfaces, *Bol. Soc. Brasil. Mat.* (*N.S.*) **20** (1990), 1–24.

[135] L. Marin, Dynamical bounds for Sturmian Schrödinger operators, *Rev. Math. Phys.* **22** (2010), 859–879.

[136] H. McCluskey, A. Manning, Hausdorff dimension for horseshoes, *Ergod. Th. & Dynam. Sys.* **3** (1983), 251–260.

[137] A. Manning, A relation between Lyapunov exponents, Hausdorff dimension and entropy, *Ergodic Theory Dynam. Systems* **1** (1981), 451–459.

[138] C. Moreira, E. Morales, J. Rivera-Letelier, On the topology of arithmetic sums of regular Cantor sets, *Nonlinearity* **13** (2000), 2077–2087.

[139] R.G. Mosier, Root neighborhoods of a polynomial, *Math. Comp.* **47** (1986), 265–273.

[140] S. Newhouse, Non-density of Axiom A(a) on S^2, *Proc. A.M.S. Symp. Pure Math.* **14** (1970), 191–202.

[141] S. Ostlund, R. Pandit, D. Rand, H. Schellnhuber, E. Siggia, One-dimensional Schrödinger equation with an almost periodic potential, *Phys. Rev. Lett.* **50** (1983), 1873–1877.

[142] S. Ostlund, R. Pandit, Renormalization-group analysis of the discrete quasiperiodic Schrödinger equation, *Phys. Rev. B* **29** (1984), 1394–1414.

[143] W. Parry, M. Pollicott, Zeta functions and the periodic orbit structure of hyperbolic dynamics, *Astérisque* **187-188** (1990), 268 pp.

[144] L. Pastur, Spectral properties of disordered systems in the one-body approximation, *Commun. Math. Phys.* **75** (1980), 179–196.

[145] J. Palis, F. Takens, *Hyperbolicity and Sensitive Chaotic Dynamics at Homoclinic Bifurcations*, Cambridge University Press, Cambridge, 1993.

[146] J. Palis, M. Viana, High dimension diffeomorphisms displaying infinitely many periodic attractors, *Ann. of Math.* **140** (1994), 207–250.

[147] B.N. Parlett, *The Symmetric Eigenvalue Problem*, SIAM Classics edition, SIAM, Philadelphia, 1998.

[148] M. Paul, Minimal symbolic flows having minimal block growth, *Math. Systems Theory* **8** (1975), 309–315.

[149] Ya. Pesin, *Dimension Theory in Dynamical Systems: Contemporary Views and Applications*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1997.

[150] M. Pollicott, Analyticity of dimensions for hyperbolic surface diffeomorphisms, to appear in *Proc. Amer. Math. Soc.*

[151] C.E. Puelz, *Improved Spectral Calculations for Discrete Schrödinger Operators*, Masters Thesis, Rice University, 2013.

[152] C. Pugh, M. Shub, A. Wilkinson, Hölder foliations, *Duke Math. J.* **86** (1997), 517–546.

[153] L. Raymond, A constructive gap labelling for the discrete Schrödinger operator on a quasiperiodic chain, preprint (1997).

[154] P. Ribeiro, P. Milman, R. Mosseri, Aperiodic quantum random walks, *Phys. Rev. Lett.* **93** (2004), 190503, 4 pp.

[155] A. Romanelli, The Fibonacci quantum walk and its classical trace map, *Phys. A.* **388** (2009), 3985–3990.

[156] J. Roberts, Escaping orbits in trace maps, *Phys. A* **228** (1996), 295–325.

[157] J. Roberts, M. Baake, Trace maps as 3D reversible dynamical systems with an invariant, *J. Statist. Phys.* **74** (1994), 829–888.

[158] D. Ruelle, *Thermodynamic Formalism. The Mathematical Structures of Classical Equilibrium Statistical Mechanics*, Encyclopedia of Mathematics and its Applications **5**, Addison-Wesley Publishing Co., Reading, 1978.

[159] O. Sarig, Lecture Notes on Thermodynamic Formalism for Topological Markov Shifts, preprint.

[160] W. Schlag, On the integrated density of states for Schrödinger operators on $\mathbb{Z}^2$ with quasi periodic potential, *Commun. Math. Phys.* **223** (2001), 47–65.

[161] B. Simon, Kotani theory for one dimensional stochastic Jacobi matrices, *Commun. Math. Phys.* **89** (1983), 227–234.

[162] B. Simon, M. Taylor, Harmonic analysis on $\mathrm{SL}(2,\mathbb{R})$ and smoothness of the density of states in the one-dimensional Anderson model, *Commun. Math. Phys.* **101** (1985), 1–19.

[163] C. Sire, Electronic spectrum of a 2D quasi-crystal related to the octagonal quasiperiodic tiling, *Europhys. Lett.* **10** (1989), 483–488.

[164] A. Sütő, The spectrum of a quasiperiodic Schrödinger operator, *Commun. Math. Phys.* **111** (1987), 409–415.

[165] A. Sütő, Singular continuous spectrum on a Cantor set of zero Lebesgue measure for the Fibonacci Hamiltonian, *J. Stat. Phys.* **56** (1989), 525–531.

[166] A. Sütő, Schrödinger difference equation with deterministic ergodic potentials, in *Beyond Quasicrystals* (*Les Houches,* 1994), Springer, Berlin (1995), 481–549.

[167] B. Sutherland, Simple system with quasiperiodic dynamics: a spin in a magnetic field, *Phys. Rev. Lett.* **57** (1986), 770–773.

[168] G. Teschl, *Jacobi Operators and Completely Integrable Nonlinear Lattices*, American Mathematical Society, Providence, R.I., 1999.

[169] A. Volberg, On the dimension of harmonic measure of Cantor repellers, *Michigan Math. J.* **40** (1993), 239–258.

[170] P. Walters, *An Introduction to Ergodic Theory*, Graduate Texts in Mathematics **79**, Springer, 1981.

[171] P. Walters, A variational principle for the pressure of continuous transformations, *Amer. J. Math.* **97** (1975), 937–971.

[172] F. Wijnands, Energy spectra for one-dimensional quasiperiodic potentials: bandwidth, scaling, mapping and relation with local isomorphism, *J. Phys. A* **22** (1989), 3267–3282.

[173] J.H. Wilkinson, The perfidious polynomial, in *Studies in Numerical Analysis*, G.H. Golub, ed., Mathematical Association of America, (1984), 1–28.

[174] J.H. Wilkinson, C. Reinsch, *Handbook for Automatic Computation, II: Linear Algebra*, Springer-Verlag, Berlin, 1971.

[175] W. Yessen, On the energy spectrum of 1D quantum Ising quasicrystal, to appear in *Ann. H. Poincaré*.

[176] W. Yessen, Properties of 1D classical and Quantum Ising quasicrystals: rigorous results, *Ann. Henri Poincaré* **15** (2014), 793–828.

[177] W. Yessen, Spectral analysis of tridiagonal Fibonacci Hamiltonians, *J. Spectr. Theory*, **3** (2013), 101–128.

[178] J.-C. Yoccoz, Some questions and remarks about $\mathrm{SL}(2,\mathbb{R})$ cocycles, *Modern Dynamical Systems and Applications*, 447–458, Cambridge Univ. Press, Cambridge, 2004.

[179] Z. Zhang, Resolvent set of Schrödinger operators and uniform hyperbolicity, preprint (arXiv:1305.4226).

David Damanik
Department of Mathematics
Rice University
Houston, TX 77005, USA
e-mail: damanik@rice.edu

Mark Embree
Department of Mathematics
Virginia Tech
225 Stanger Street 0123
Blacksburg, Virginia 24061, USA
e-mail: embree@rice.edu

Anton Gorodetski
Department of Mathematics
University of California
Irvine CA 92697, USA
e-mail: asgor@math.uci.edu

Progress in Mathematics, Vol. 309, 371–403
© 2015 Springer Basel

Additive Properties of Sets and Substitutive Dynamics

Svetlana Puzynina and Luca Q. Zamboni

Abstract. In the chapter we discuss a new connection between central sets and the strong coincidence conjecture for fixed points of irreducible primitive substitutions of Pisot type. Central sets, first introduced by Furstenberg using notions from topological dynamics, constitute a special class of subsets of $\mathbb{N}$ possessing strong combinatorial properties: Each central set contains arbitrarily long arithmetic progressions, and solutions to all partition regular systems of homogeneous linear equations. We give an equivalent reformulation of the strong coincidence condition in terms of central sets and minimal idempotent ultrafilters in the Stone–Čech compactification $\beta\mathbb{N}$. This provides a new arithmetical approach to an outstanding conjecture in tiling theory, the Pisot substitution conjecture. Using various families of uniformly recurrent words, including Sturmian words, the Thue–Morse word and fixed points of weak mixing substitutions, we generate an assortment of central sets which reflect the rich combinatorial structure of the underlying words. One crucial additive property of central sets is that each central set contains all finite sums of distinct terms for some infinite increasing sequence of natural numbers, i.e. is an IP-set. By a celebrated result of $\mathbb{N}$. Hindman, the collection of all IP-sets is partition regular, i.e., if A is an IP-set then for any finite partition of A, one cell of the partition is an IP-set. We introduce an hierarchy of additive combinatorial properties for subsets of $\mathbb{N}$ and study them in terms of partition regularity. The results introduced in the chapter rely on interactions between different areas of mathematics: They include the general theory of combinatorics on words, numeration systems, tilings, topological dynamics and the algebraic/topological properties of Stone–Čech compactification of $\mathbb{N}$.

Mathematics Subject Classification (2010). 37B10, 68R15, 05D10, 11A63.

Keywords. Substitutions, central sets, IP-sets, Pisot conjecture, partition regularity, additive combinatorics, abstract numeration systems.

1. Introduction

An important open problem in the theory of substitutions is the so-called *strong coincidence conjecture*: It states that each pair of fixed points x and y of an irreducible primitive substitution of Pisot type are *strongly coincident*: There exist a letter a and a pair of Abelian equivalent words s, t, such that sa is a prefix of x and ta is a prefix of y. This combinatorial condition, originally due to P. Arnoux and S. Itô, is an extension of a similar condition considered by F.M. Dekking in [18] in the case of uniform substitutions. In this case Dekking proves that the condition is satisfied by the "pure base" of the substitution if and only if the associated substitutive subshift has *pure discrete spectrum*, i.e., is metrically isomorphic with translation on a compact Abelian group. The strong coincidence conjecture has been verified for irreducible primitive substitutions of Pisot type on a binary alphabet in [3] and is otherwise still open.

The strong coincidence conjecture is linked to diffraction properties of one-dimensional atomic arrangements in the following way. It is shown in [22] and [29] that an atomic arrangement determined by a substitution has pure point diffraction spectrum (i.e., is a perfect quasicrystal) if and only if the tiling system associated with the substitution has pure discrete dynamical spectrum. In a pair of seminal papers ([35, 36]), G. Rauzy established a link between pure discreteness of the dynamical spectrum and the irreducible Pisot property for substitutions. The *Pisot substitution conjecture* asserts that the dynamical spectrum of the tiling system associated with an irreducible Pisot substitution has pure discrete dynamical spectrum. For the latter to hold, it is necessary, and conjecturally sufficient, for the substitution to satisfy the strong coincidence condition.

In this chapter we describe a connection between the strong coincidence conjecture and *central sets*, originally introduced by Furstenberg in [24]. Central sets are subsets of the natural numbers $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$ which are known to have substantial combinatorial structure of additive nature. For example, any central set contains arbitrarily long arithmetic progressions, and solutions to all partition regular systems of homogeneous linear equations (see for example [10]). Furstenberg originally defined them using notions from topological dynamics.

One crucial additive property of central sets is that each central sets contains all finite sums of distinct terms for some infinite sequence of natural numbers $x_0 < x_1 < x_2 \cdots$. More precisely, let $\mathrm{Fin}(\mathbb{N})$ denote the set of all non-empty finite subsets of $\mathbb{N}$. A subset A of $\mathbb{N}$ is called an *IP-set* if A contains $\{\sum_{n\in F} x_n \,|\, F \in \mathrm{Fin}(\mathbb{N})\}$ for some infinite sequence of natural numbers $x_0 < x_1 < x_2 \cdots$. Every central set is an IP-set but not conversely.

By a celebrated result of N. Hindman [25], given any finite partition of $\mathbb{N}$, at least one element of the partition is an IP-set. More generally Hindman shows that given any finite partition of an IP-set, at least one element of the partition is again an IP-set. In other words the property of being an IP-set is *partition regular*, i.e., cannot be destroyed via a finite partitioning. Other examples of partition regularity

are given by the pigeonhole principle, sets having positive upper density, and sets having arbitrarily long arithmetic progressions (Van der Waerden's theorem).

Both IP-sets and central sets may be characterized in terms of the algebro/topological properties of the Stone–Čech compactification $\beta\mathbb{N}$. We regard $\beta\mathbb{N}$ as the collection of all ultrafilters on $\mathbb{N}$. There is a natural extension of the operation of addition $+$ on $\mathbb{N}$ to $\beta\mathbb{N}$ making $\beta\mathbb{N}$ a compact *left-topological semigroup*. Via a celebrated result of Ellis [23], $\beta\mathbb{N}$ contains *idempotents*, i.e., ultrafilters $p \in \beta\mathbb{N}$ satisfying $p + p = p$. A striking result due to Hindman links IP-sets and idempotents in $\beta\mathbb{N}$: A subset $A \subseteq \mathbb{N}$ is an IP-set if and only if $A \in p$ for some idempotent $p \in \beta\mathbb{N}$ (see Theorem 5.12 in [27]).

In [9], Bergelson and Hindman showed that central sets too may alternatively be defined in terms of a special class of ultrafilters, called *minimal idempotents*. Every compact Hausdorff left-topological semigroup $\mathcal{S}$ admits a smallest two-sided ideal $K(\mathcal{S})$ which is at the same time the union of all minimal right ideals of $\mathcal{S}$ and the union of all minimal left ideals of $\mathcal{S}$ (see for instance [27]). It is readily verified that the intersection of any minimal left ideal with any minimal right ideal is a group. In particular, there are idempotents in $K(\mathcal{S})$. Such idempotents are called minimal and their elements are called central sets, i.e., $A \subset \mathbb{N}$ is a *central set* if it is a member of some minimal idempotent in $\beta\mathbb{N}$. It follows from this that the collection of all central sets is also partition regular.

The question of determining whether a given subset $A \subseteq \mathbb{N}$ is an IP-set or a central set is typically quite difficult, even if for every A, either A or its complement is an IP-set (resp. central set). Although as we saw both IP-sets and central sets are characterized as belonging to a certain class of idempotent ultrafilters, the question of belonging or not to a given (non-principal) ultrafilter is generally equally mysterious. An equivalent word combinatorial reformulation of this question is as follows: Given a binary word $\omega = \omega_0\omega_1\omega_2\cdots \in \{0,1\}^\infty$, put $\omega\big|_0 = \{n \in \mathbb{N} \,|\, \omega_n = 0\}$ and $\omega\big|_1 = \{n \in \mathbb{N} \,|\, \omega_n = 1\}$. The question is then to determine whether the set $\omega\big|_0$ or $\omega\big|_1$ is an IP-set or central set. Of course in general, this reformulation is as difficult as the original question. However, should the word ω be characterized by some rich combinatorial properties, or be generated by some "simple" combinatorial or geometric algorithm (such as a substitution rule, a finite state automaton, a Toeplitz rule ...) or arise as a natural coding of a reasonably simple symbolic dynamical system, then the underlying rigid combinatorial structure of the word may provide insight to our previous question. Furthermore, such families of words may be used to obtain simple constructions of central sets having additional nice properties inherited from the rich underlying combinatorial structure. One of our objectives in this chapter is to illustrate this latter point.

An ultrafilter may be thought of as a $\{0,1\}$-valued finitely additive probability measure defined on all subsets of $\mathbb{N}$. This notion of measure induces a notion of convergence (p-$\lim_n$) for sequences indexed by $\mathbb{N}$, which we regard as a mapping from words to words. Using this key notion of convergence, the second

author together with M. Barge obtained the following reformulation of the strong coincidence conjecture in terms of both central sets IP-sets:

Theorem 1 ([7]). *Let τ be an irreducible primitive substitution of Pisot type. Then for any pair of fixed points x and y of τ the following are equivalent:*

1. *x and y are strongly coincident.*
2. *There exists a minimal idempotent $p \in \beta\mathbb{N}$ such that $y = p\text{-}\lim_n T^n(x)$ where T denotes the shift map.*
3. *For any prefix u of y, the set of occurrences of u in x is a central set.*
4. *For any prefix u of y, the set of occurrences of u in x is an IP-set.*

Since IP-sets may be defined arithmetically in terms of finite sums of distinct terms of infinite sequences of natural numbers, the above theorem provides an arithmetical approach to solving the strong coincidence conjecture. To this end, certain numeration systems first introduced by J.-M. Dumont and A. Thomas in [20, 21] may be useful for the conjecture.

The chapter is organized as follows: In §2 we introduce a hierarchy of additive combinatorial properties for subsets of $\mathbb{N}$. We study each one in terms of partition regularity. At first, we take a very simplistic approach using only basic methods of combinatorics on words. In §3 we review some fundamental features of the algebro/topological properties of the Stone–Čech compactification $\beta\mathbb{N}$ from which much deeper properties of IP-sets and central sets are derived. In §4 we use various infinite words (including Sturmian words and primitive morphic words) to produce IP-sets and central sets having additional rich combinatorial properties inherited from the underlying words. In §5 we describe the connection between central sets and the strong coincidence condition from primitive substitutions particularly for those of Pisot type. In §6 we discuss the so-called Dumont–Thomas numeration systems defined by substitutions. Finally in §7 we give a few open problems.

2. Additive properties of subsets of $\mathbb{N}$

2.1. Definitions and examples

Let $\mathbb{N} = \{0, 1, 2, \ldots\}$ denote the set of natural numbers and $\mathbb{N}^+ = \{1, 2, 3, \ldots\}$ the set of positive integers. We consider a hierarchy of additive properties satisfied by some subsets on the natural numbers $\mathbb{N}$. Let $x_1, x_2, \ldots, x_k$ be a k-term sequence of natural numbers. We do not assume that the x_i are pairwise distinct. We set

$$FS(x_t)_{t=1}^{k} = \left\{ \sum_{t \in F} x_t \mid F \subset \{1, 2, \ldots, k\} \right\}.$$

Given an infinite sequence $x_1, x_2, x_3, \ldots$ of natural numbers, we set

$$FS_{\leq k}(x_t)_{t=1}^{\infty} = \left\{ \sum_{t \in F} x_t \mid F \subset \mathbb{N}^+, \#F \leq k \right\}$$

and

$$FS(x_t)_{t=1}^{\infty} = \left\{ \sum_{t \in F} x_t \,|\, F \in \mathrm{Fin}(\mathbb{N}^+) \right\}$$

where $\mathrm{Fin}(\mathbb{N}^+)$ denotes the set of all finite subsets of $\mathbb{N}^+$.

Definition 2.1. Let $k \in \mathbb{N}^+$ and $A \subseteq \mathbb{N}$. We say A is *k-summable* if A contains $FS(x_t)_{t=1}^{k}$ for some k-term sequence of natural numbers $(x_t)_{t=1}^{k}$. We say A is *k^∞-summable* if A contains $FS_{\leq k}(x_t)_{t=1}^{\infty}$ for some infinite sequence of natural numbers $(x_t)_{t=1}^{\infty}$.

Definition 2.2. Let $A \subseteq \mathbb{N}$. We say A is *finite FS-big* if A is k-summable for each positive integer k. We say A is *infinite FS-big* if A is k^∞-summable for each positive integer k.

Definition 2.3. Let $A \subseteq \mathbb{N}$. We say A is an *IP-set* if A contains $FS(x_t)_{t=1}^{\infty}$ for some infinite sequence of natural numbers $(x_t)_{t=1}^{\infty}$.

Let Σ_k ($\Sigma_k^\infty, \Sigma, \Sigma^\infty$, and $\mathcal{IP}$ respectively) denote the set of all k-summable (k^∞-summable, finite FS-big, infinite FS-big, and IP) sets. Then $\Sigma = \bigcap_{k \geq 1} \Sigma_k$ and $\Sigma^\infty = \bigcap_{k \geq 1} \Sigma_k^\infty$. Moreover we have the following inclusions $\Sigma_{k+1} \subseteq \Sigma_k$, $\Sigma_{k+1}^\infty \subseteq \Sigma_k^\infty$, $\Sigma_k^\infty \subseteq \Sigma_k$ and

$$\mathcal{IP} \subseteq \Sigma^\infty \subseteq \Sigma.$$

As we shall see, each of the above inclusions is strict. For $\mathcal{S} \in \{\Sigma, \Sigma^\infty, \mathcal{IP}\}$, we denote by $\mathcal{S}^*$ the set of all subsets A of $\mathbb{N}$ such that $A \cap B \neq \emptyset$ for all $B \in \mathcal{S}$.

The set of even numbers $\{0, 2, 4, \ldots\}$ is perhaps the simplest example of an IP-set as it is closed under addition. In contrast the set of odd numbers $\{1, 3, 5, 7, \ldots\}$ is not 2-summable. It is readily verified that for each $k > 2$, the set $\{n \in \mathbb{N} \,|\, n \not\equiv 0 \bmod k\} \in \Sigma_{k-1}^\infty \setminus \Sigma_k$ (see [14]). We now show how to generate interesting and non-trivial examples of IP-sets starting from combinatorially rich infinite words, such as those studied in the field of combinatorics on words. Let $\mathbb{A}$ denote a finite non-empty set (called the alphabet), we denote by $\mathbb{A}^*$, $\mathbb{A}^{\mathbb{N}}$ and $\mathbb{A}^+$ respectively the set of finite words, the set of infinite words, and the set of non-empty finite words over the alphabet $\mathbb{A}$. Let $\omega = \omega_0\omega_1\omega_2 \cdots \in \mathbb{A}^{\mathbb{N}}$ be an infinite word on the alphabet $\mathbb{A}$. For each finite word $u \in \mathbb{A}^+$ we set

$$\omega\big|_u = \{n \in \mathbb{N} \,|\, \omega_n\omega_{n+1} \cdots \omega_{n+|u|-1} = u\}.$$

In other words, $\omega\big|_u$ denotes the set of all occurrences of u in ω. If $\omega\big|_u$ is non-empty, u is called a *factor* of ω; if $0 \in \omega\big|_u$, then u is a *prefix* of ω. We say ω is *recurrent* if for each prefix u of ω the set $\omega\big|_u$ is infinite. We say that ω is *uniformly recurrent* if for each prefix u of ω the set $\omega\big|_u$ is syndetic, i.e., of bounded gap.

Proposition 2.4. *Let $\omega = \omega_0\omega_1\omega_2 \cdots \in \mathbb{A}^{\mathbb{N}}$ be recurrent and set $a = \omega_0$. Then $\omega\big|_a$ is an IP-set.*

Proof. Set $x_0 = 0$ and $B_0 = a$. Since ω is recurrent, the prefix B_0 of ω occurs again in ω at some position $x_1 > x_0$. Let B_1 denote the prefix of ω of length $x_1 + 1$. Since ω is recurrent, B_1 must occur at some position $x_2 > x_1$. Since B_1 begins and ends in B_0 we have that $FS(x_t)_{t=0}^2 \in \omega\big|_a$. Let B_2 be the prefix of ω of length $x_2 + x_1 + 1$. By recurrence, B_2 must occur in ω at some position $x_3 > x_2 + x_1$. Since B_2 begins and ends in B_1 we have that $FS(x_t)_{t=0}^3 \in \omega\big|_a$. Let B_3 be the prefix of ω of length $x_3 + x_2 + x_1 + 1$ (see Figure 1). Continuing in this way wee construct an infinite sequence of natural numbers $0 = x_0 < x_1 < x_2 < \cdots$ such that $FS(x_t)_{t=0}^{\infty} \in \omega\big|_a$. □

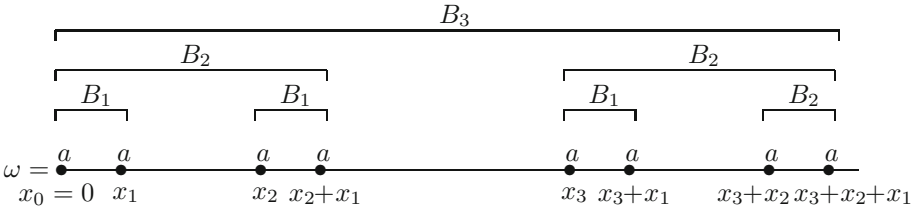

FIGURE 1. Proof of Proposition 2.4.

Thus every recurrent infinite word defines an IP-set simply by taking the set of occurrences of its first symbol.

Now we are going to consider several examples defined by substitutional rules. A *substitution* τ on an alphabet $\mathbb{A}$ is a mapping $\tau : \mathbb{A} \to \mathbb{A}^+$. The mapping τ extends by concatenation to maps (also denoted τ) $\mathbb{A}^* \to \mathbb{A}^*$ and $\mathbb{A}^{\mathbb{N}} \to \mathbb{A}^{\mathbb{N}}$. A word $\omega \in \mathbb{A}^{\mathbb{N}}$ is called a *fixed point* of τ if $\tau(\omega) = \omega$.

Example 1: The Thue–Morse word. Consider the *Thue–Morse infinite word*

$$\mathbb{T} = t_0 t_1 t_2 t_3 \cdots = 01101001100101 \cdots$$

where t_n is defined as the sum modulo 2 of the digits in the binary expansion of n. The Thue–Morse word is 2-automatic [2]: In fact t_n is computed by feeding the binary expansion of n in the deterministic finite automata depicted in Figure 2. Starting from the initial state labeled 0, we read the binary expansion of n starting from the most significant digit. Then t_n is the corresponding terminal state. For example, the binary representation of 13 is 1101 and the path 1101 starting at 0 terminates at vertex 1. Whence $t_{13} = 1$.

The origins of $\mathbb{T}$ go back to the beginning of the last century with the works of the Norwegian mathematician Axel Thue [38]. Thue noted that every binary word of length four contains a square, that is two consecutive equal blocks XX. He then asked whether it was possible to find an infinite word on 3 distinct symbols which avoided squares. He also asked whether there exists an infinite binary word without cubes, that is with no three consecutive equal blocks. Thue showed that in each case the answer is positive and constructed this very special infinite word $\mathbb{T}$ to produce the desired words. In fact the word $\mathbb{T}$ contains no fractional power

greater than 2, i.e., contains no word of the form XXX' where X' is a prefix of X. Thue's work originally appeared in an obscure Norwegian journal and for many years remained largely unknown and unappreciated. This word, originally defined by Thue to study combinatorial properties of words, was rediscovered in 1921 by Morse [30] in connection with differential geometry.

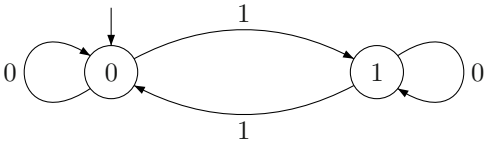

FIGURE 2. The Thue–Morse automaton.

It follows from Proposition 2.4 that $\mathbb{T}\big|_0$ is an IP-set. This could have also been proved directly from the base 2 numeration system (see Theorem 3.1 in [14]). For each positive integer n we will denote the binary expansion of n by $[n]_2$, i.e., if $n = r_k 2^k + r_{k-1}2^{k-1} + \cdots + r_0 2^0$ with $r_k = 1$ and $r_i \in \{0,1\}$ we write $[n]_2 = r_k r_{k-1} \cdots r_0$. We define the *support* of n, denote $\mathrm{supp}(n)$ by $\mathrm{supp}(n) = \{i \in \{0,1,\ldots,k\} \,|\, r_i = 1\}$. For instance, $\mathrm{supp}(19) = \{0,1,4\}$. Thus

$$n \in \mathbb{T}\big|_0 \Leftrightarrow t_n = 0 \Leftrightarrow \#\mathrm{supp}(n) \text{ is even.}$$

Put $x_n = 2^{2n+1} + 2^{2n}$. Then clearly $FS(x_n)_{n=0}^{\infty} \subset \mathbb{T}\big|_0$. Whence $\mathbb{T}\big|_0$ is an IP-set. We shall see later that $\mathbb{T}\big|_1$ is not an IP-set from which it follows that $\mathbb{T}\big|_0 \in \mathcal{IP}^*$.

Example 2: The Fibonacci word. As another example we consider the *Fibonacci* infinite word

$$\mathbf{f} = f_0 f_1 f_2 \cdots = 0100101001001010010100100101001010010\cdots$$

fixed by the morphism $0 \mapsto 01$, $1 \mapsto 0$. Equivalently, this word can be defined via recurrence relations. Namely, set $\mathbf{f}_0 = 0$, $\mathbf{f}_1 = 01$, and $\mathbf{f}_{n+1} = \mathbf{f}_n \mathbf{f}_{n-1}$. Then the Fibonacci word can be defined as a limit: $\mathbf{f} = \lim_{n\to\infty} \mathbf{f}_n$. We remark that the sequence of lengths $|\mathbf{f}_n|$ is given by the sequence of Fibonacci numbers.

It follows from Proposition 2.4 that

$$\mathbf{f}\big|_0 = \{0,2,3,5,7,8,10,11,13,15,16,\ldots\}$$

is an IP-set. As in the previous example, this could have been proved directly from the underlying numeration system. Let $F_0 = 1, F_1 = 2, F_2 = 3, \ldots$ be the sequence of Fibonacci numbers. It is well known that each positive integer n has one or more representations when expressed as a sum of distinct Fibonacci numbers. One way of obtaining such a representation is by applying the 'greedy algorithm.' This gives rise to a representation of n of the form $n = \sum_{i=0}^{k} t_i F_i$ with $t_i \in \{0,1\}$ and with $t_{i+1}t_i \neq 11$ for each $0 \leq i \leq k-1$, called the *Zeckendorff representation* [40]. We shall write $\mathcal{Z}(n) = t_k t_{k-1} \cdots t_0$. For example, applying the greedy algorithm to $n = 50$ we obtain $50 = 34 + 13 + 3 = F_7 + F_5 + F_2$ which gives rise to the representation $\mathcal{Z}(50) = 10100100$. It follows that $\mathcal{Z}(F_n) = 10^n$. It is also well

known that for each $n \in \mathbb{N}$, $f_n = 0$ whenever $\mathcal{Z}(n)$ ends in 0 and $f_n = 1$ whenever $\mathcal{Z}(n)$ ends in 1. Thus

$$\mathbf{f}\big|_0 = \{n \in \mathbb{N} \mid \mathcal{Z}(n) \text{ ends in } 0\} \qquad \text{and} \qquad \mathbf{f}\big|_1 = \{n \in \mathbb{N} \mid \mathcal{Z}(n) \text{ ends in } 1\}.$$

Consider sequence $(x_n)_{n\in\mathbb{N}}$ given by $x_n = F_{2n+1}$. It is readily verified that for each $A \in \mathrm{Fin}(\mathbb{N})$, the Zeckendorff representation of $\sum_{n\in A} x_n$ ends in 10^{2m+1} where $m = \min(A)$. Thus not only is each x_n in $\mathbf{f}\big|_0$, but so is every sum of the form $\sum_{n\in A} x_n$ with $A \in \mathrm{Fin}(\mathbb{N})$. Thus we have shown that $\mathbf{f}\big|_0$ is an IP-set. We shall see later that $\mathbf{f}\big|_1$ is not an IP-set, whence $\mathbf{f}\big|_0 \in \mathcal{IP}^*$.

Example 3: The Sierpinski word. As a last example, consider the Sierpinski word

$$\mathbf{s} = 0101110101111111101011101\cdots$$

fixed by the morphism $0 \mapsto 010$ $1 \mapsto 111$. It follows from Proposition 2.4 that $\mathbf{s}\big|_0$ is an IP-set. Unlike in the previous two examples, this IP-set is not syndetic. On the other hand, it is easy to see that $\mathbf{s}_1$ is also an IP-set. This set contains in particular all numbers between 3^i and $2 \cdot 3^i - 1$ and hence all numbers of the form $\Sigma_{i\in A} 3^i$, where $A \in \mathrm{Fin}(\mathbb{N})$. So, $\mathbf{s}_1$ is IP-set by definition, the corresponding sequence is $x_i = 3^i$. So, the complement of $\mathbf{s}_0$ is also an IP-set, whence $\mathbf{s}\big|_0 \notin \mathcal{IP}^*$.

2.2. Partition regularity

A fundamental result in Ramsey theory, originally due to Issai Schur [37], states that given a finite partition of the positive integers $\mathbb{N}^+$, one cell of the partition contains two points x, y and their sum $x + y$. In other words, one cell of the partition is 2-summable. While in Schur's theorem it is not assumed that x and y are distinct, the theorem would no longer be true if we insisted that $x = y$. In fact, for $i \in \{0, 1\}$, let A_i denote the set of all positive integers n such that writing $n = 2^j m$ with m odd we have $j \equiv i \pmod 2$. Then clearly no A_i contains both x and $2x$. For each positive integer n, let $\omega_n = i \in \{0, 1\}$ if $n \in A_i$. Then the infinite word $\omega = \omega_1\omega_2\omega_3\cdots = 01000101\cdots$ is the well-known *period doubling word* defined as the fixed point of the morphism $0 \mapsto 01$, $1 \mapsto 00$. (see [1]).

An extension of Schur's Theorem, which we will call the *finite Finite Sums Theorem* states that whenever $\mathbb{N}^+$ is finitely partitioned, then for each positive integer k one cell of the partition is k-summable. The finite Finite Sums Theorem is a straightforward consequence of Rado's Theorem [33]. It follows from this and the pigeon hole principle that given any finite partition of $\mathbb{N}^+$, one cell of the partition is finite FS-big. In 1974, N. Hindman proved the following infinite extension of the finite Finite Sums Theorem:

Theorem 2.5 (N. Hindman, [25]). *Given any finite partition of* $\mathbb{N}^+$, *one cell of the partition is an IP-set, i.e., contains all finite sums of distinct terms of some infinite sequence* $(x_t)_{t=1}^{\infty}$.

So for any choice of $\mathcal{S} \in \{\Sigma, \Sigma^\infty, \mathcal{IP}\}$, and for each finite partition of $\mathbb{N}^+$, one cell of the partition is in $\mathcal{S}$. But one may ask something stronger:

Question 2.6. *For $\mathcal{S} \in \{\Sigma, \Sigma^\infty, \mathcal{IP}\}$, and $A \in \mathcal{S}$, is it the case that for each finite partition of A, one cell of the partition belongs to $\mathcal{S}$?*

We recall that a collection of sets $\mathcal{S}$ is said to be *partition regular* if for each $A \in \mathcal{S}$, whenever A is finitely partitioned, one cell of the partition is in $\mathcal{S}$. Examples of partition regular families are the collection of all infinite sets, or the collection of all subset of $\mathbb{N}$ having positive upper density, or the collection of all subsets of $\mathbb{N}$ having arbitrarily long arithmetic progressions. (This latter fact is an almost immediate consequence of van der Waerden's Theorem [39]. In fact, assume $A \subseteq \mathbb{N}$ contains arbitrarily long arithmetic progressions. Let $k, r \in \mathbb{N}$, and let $A = \bigcup_{i=1}^r C_i$. By van der Waerden's Theorem pick n such that whenever $\{1, 2, \ldots, n\}$ is partitioned into r classes, one class contains a length k arithmetic progression. Pick a and d in $\mathbb{N}$ such that $\{a + d, a + 2d, \ldots, a + nd\} \subseteq A$. For $i \in \{1, \ldots, r\}$ let $B_i = \{t \in \{1, 2, \ldots, n\} \mid a + td \in C_i\}$. Pick i, b, and c such that $\{b+c, b+2c, \ldots, b+kc\} \subseteq B_i$. Then $\{a+bd+cd, a+bd+2cd, \ldots, a+bd+kcd\} \subseteq C_i$.)

N. Hindman proved that $\mathcal{IP}$ is partition regular:

Theorem 2.7 ([27]). *The collection of all IP-sets is partition regular.*

It is easy to see that the sets Σ_k and Σ_k^∞ are not partition regular. For instance, the set $A = \{n \in \mathbb{N}^+ \mid n \equiv 1, 2 \pmod 3\}$ is Σ_2^∞. On the other hand, neither set $\{n \in \mathbb{N}^+ \mid n \equiv 1 \pmod 3\}$, $\{n \in \mathbb{N}^+ \mid n \equiv 2 \pmod 3\}$ is in Σ_2. In [14], the authors together with M. Bucci and N. Hindman proved that Σ is partition regular.

Theorem 2.8 ([14]). *The collection of all finite FS-big sets is partition regular.*

Before we proceed to the proof we observe that we had some choices to make when we defined k-summable. That is, we could have defined A to be k-summable$_1$ if there is a sequence $\langle x_t \rangle_{t=1}^k$ such that $FS(\langle x_t \rangle_{t=1}^k) \subseteq A$; we could have defined A to be k-summable$_2$ if there is an increasing sequence $\langle x_t \rangle_{t=1}^k$ such that $FS(\langle x_t \rangle_{t=1}^k) \subseteq A$; and we could have defined A to be k-summable$_3$ if there is a sequence $\langle x_t \rangle_{t=1}^k$ satisfying uniqueness of finite sums such that $FS(\langle x_t \rangle_{t=1}^k) \subseteq A$. These notions are progressively strictly stronger. For example if $k > 1$, $\{1, 2, \ldots, k\}$ is k-summable$_1$ but not k-summable$_2$. And if $k > 1$, $\left\{1, 2, \ldots, \frac{k^2+k}{2}\right\}$ is k-summable$_2$ but not k-summable$_3$. However, for the notion of finite FS-big subsets of $\mathbb{N}$, it does not matter which choice was made for k-summable. The reason is that for each k there is some m such that if A is an m-summable$_1$ subset of $\mathbb{N}$, then A is k-summable$_2$. Similarly, for each k there is some m such that if A is an m-summable$_2$ subset of $\mathbb{N}$, then A is k-summable$_3$.

The main key for proving that finite FS-big sets are partition regular is the finite Finite Unions Theorem. The proof that we will present uses a standard compactness argument and the infinite Finite Unions Theorem.

Theorem 2.9 (Infinite Finite Unions Theorem). *Let $r \in \mathbb{N}$. If $\mathrm{Fin}(\mathbb{N}) = \bigcup_{i=1}^r \mathcal{F}_i$, then there exist $i \in \{1, 2, \ldots, r\}$ and a sequence $\langle F_t \rangle_{t=1}^\infty$ in $\mathrm{Fin}(\mathbb{N})$ such that for each $t \in \mathbb{N}$, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\mathbb{N})$, $\bigcup_{t \in H} F_t \in \mathcal{F}_i$.*

Proof. This is actually stated in [25]. A much easier proof is in [27, Cor. 5.17]. It is an immediate corollary of Theorem 2.7, because, given any sequence $\langle x_t\rangle_{t=1}^{\infty}$ in $\mathbb{N}$ one may choose a sequence $\langle F_n\rangle_{n=1}^{\infty}$ in $\mathrm{Fin}(\mathbb{N})$ such that for each $n \in \mathbb{N}$, $\max F_n < \min F_{n+1}$ and for each n and l in $\mathbb{N}$, if $2^l \leq \sum_{t\in F_n} x_t$, then 2^{l+1} divides $\sum_{t\in F_{n+1}} x_t$. (That is, the maximum of the binary support of $\sum_{t\in F_n} x_t$ is less than the minimum of the binary support of $\sum_{t\in F_{n+1}} x_t$.) □

Theorem 2.10 (Finite Finite Unions Theorem). *Let $r, k \in \mathbb{N}$. There is some $m \in \mathbb{N}$ such that whenever $\mathrm{Fin}(\{1,2,\ldots,m\}) = \bigcup_{i=1}^{r} \mathcal{F}_i$, there exist $i \in \{1,2,\ldots,r\}$ and a sequence $\langle F_t\rangle_{t=1}^{k}$ in $\mathrm{Fin}(\{1,2,\ldots,m\})$ such that for each $t \in \{1,2,\ldots,k-1\}$, if any, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$, $\bigcup_{t\in H} F_t \in \mathcal{F}_i$.*

Proof. Suppose not. For each $m \in \mathbb{N}$ pick a function $\psi_m : \mathrm{Fin}(\{1,2,\ldots,m\}) \to \{1,2,\ldots,r\}$ with the property that there do not exist $i \in \{1,2,\ldots,r\}$ and a sequence $\langle F_t\rangle_{t=1}^{k}$ in $\mathrm{Fin}(\{1,2,\ldots,m\})$ such that for each $t \in \{1,2,\ldots,k-1\}$, if any, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$, $\psi_m(\bigcup_{t\in H} F_t) = i$. Define $\sigma_m : \mathrm{Fin}(\mathbb{N}) \to \{1,2,\ldots,r\}$ by $\sigma_m(F) = \psi_m(F)$ if $F \subseteq \{1,2,\ldots,m\}$ and $\sigma_m(F) = 1$ otherwise.

Give $\{1,2,\ldots,r\}$ the discrete topology and let $X = \bigtimes_{F\in\mathrm{Fin}(\mathbb{N})}\{1,2,\ldots,r\}$ with the product topology. Then X is compact and $\langle\sigma_m\rangle_{m=1}^{\infty}$ is a sequence in X so pick a cluster point φ of $\langle\sigma_m\rangle_{m=1}^{\infty}$. Pick by Theorem 2.9, $i \in \{1,2,\ldots,r\}$ and a sequence $\langle F_t\rangle_{t=1}^{\infty}$ in $\mathrm{Fin}(\mathbb{N})$ such that for each $t \in \mathbb{N}$, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\mathbb{N})$, $\varphi(\bigcup_{t\in H} F_t) = i$. Let

$$U = \left\{\mu \in X \,|\, \mu(F_i) = \varphi(F_i) \text{ for all } i \in \{1,2,\ldots,k\}\right\}.$$

Then U is a neighborhood of φ in X so pick $m > \max F_k$ such that $\sigma_m \in U$. Then for each $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$, $\psi_m(\bigcup_{t\in H} F_t) = \sigma_m(\bigcup_{t\in H} F_t) = \varphi(\bigcup_{t\in H} F_t) = i$, a contradiction. □

Proof of Theorem 2.8. Suppose $A \subseteq \mathbb{N}$ is finite FS-big and $A = \bigcup_{i=1}^{r} B_i$ for some $r \in \mathbb{N}$. Let $k \in \mathbb{N}$. We shall show that there are some $i \in \{1,2,\ldots,r\}$ and some $\langle x_t\rangle_{t=1}^{k}$ satisfying uniqueness of finite sums such that $FS(\langle x_t\rangle_{t=1}^{k}) \subseteq B_i$. By the pigeon hole principle, there is thus one i which contains such a set for arbitrarily large k, and thus for all k.

By Theorem 2.10 pick $m \in \mathbb{N}$ such that whenever $\mathrm{Fin}(\{1,2,\ldots,m\}) = \bigcup_{i=1}^{r} \mathcal{F}_i$, then there exist $i \in \{1,2,\ldots,r\}$ and a sequence $\langle F_t\rangle_{t=1}^{k}$ in $\mathrm{Fin}(\{1,2,\ldots,m\})$ such that for each $t \in \{1,2,\ldots,k-1\}$, if any, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$, $\bigcup_{t\in H} F_t \in \mathcal{F}_i$. Since A is finite FS-big we may pick $\langle y_t\rangle_{t=1}^{m}$ satisfying uniqueness of finite sums with $FS(\langle y_t\rangle_{t=1}^{m}) \subseteq A$. For each $i \in \{1,2,\ldots,r\}$ let $\mathcal{F}_i = \{H \in \mathrm{Fin}(\{1,2,\ldots,m\}) \,|\, \sum_{t\in H} y_t \in B_i\}$. Pick $i \in \{1,2,\ldots,r\}$ and a sequence $\langle F_t\rangle_{t=1}^{k}$ in $\mathrm{Fin}(\{1,2,\ldots,m\})$ such that for each $t \in \{1,2,\ldots,k-1\}$, if any, $\max F_t < \min F_{t+1}$ and for each $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$, $\bigcup_{t\in H} F_t \in \mathcal{F}_i$. For $n \in \{1,2,\ldots,k\}$ let $x_n = \sum_{t\in F_n} y_t$. Then since $\max F_t < \min F_{t+1}$ when $t < k$, if $H \in \mathrm{Fin}(\{1,2,\ldots,k\})$ and $K = \bigcup_{n\in H} F_n$,

then $\sum_{n\in H} x_n = \sum_{t\in K} y_t \in B_i$. Further it is an easy exercise to show that, since $\langle y_t\rangle_{t=1}^m$ satisfies uniqueness of finite sums, so does $\langle x_t\rangle_{t=1}^k$. □

It turns out that the collection of all infinite FS-big sets is not partition regular. For this we return to the Thue–Morse word $\mathbb{T}$.

Lemma 2.11 ([14]). *There exists a partition of the set $\mathbb{T}\big|_1$ into two sets neither of which is in Σ_2^∞.*

Proof. Consider the partition $\mathbb{T}\big|_1 = A_0 \cup A_1$ defined as follows: Let A_0 be the set of all $n \in \mathbb{T}\big|_1$ such that the $\min\big(\mathrm{supp}(n)\big)$ is even, and let A_1 be the set of all $n \in \mathbb{T}\big|_1$ such that the $\min\big(\mathrm{supp}(n)\big)$ is odd. For instance, $26 = 2^4 + 2^3 + 2^1$, and hence the least nonzero digit in the binary expansion of 26 is in position 1, so $26 \in A_1$. We will show that neither A_i is in Σ_2^∞. Fix $i \in \{0, 1\}$ and suppose to the contrary that A_i is in Σ_2^∞, i.e., there is an infinite sequence $\langle x_n\rangle_{n=1}^\infty$ in A_i such that for every $n \neq m$ we have $x_n + x_m \in A_i$. Note first that for each $n > 1$ we have $\mathrm{supp}(x_n) \cap \mathrm{supp}(x_1) \neq \emptyset$. Otherwise $\#\mathrm{supp}(x_1 + x_n)$ would be even. Therefore, there exists a positive constant M such that $\min\big(\mathrm{supp}(x_n)\big) \leq M$ for each $n \in \mathbb{N}$. By the pigeon hole principle there exists a positive integer r and an infinite subsequence $x_{n_1}, x_{n_2}, \ldots$ of the sequence $\langle x_n\rangle_{n=1}^\infty$ such that $\min\big(\mathrm{supp}(x_{n_j})\big) = r$ for each $j \in \mathbb{N}$. Again by the pigeon hole principle there exists infinitely many of the x_{n_j} whose binary expansions also agree in position $r+1$. Thus there exists $n \neq m$ such that $\min\big(\mathrm{supp}(x_n)\big) = \min\big(\mathrm{supp}(x_m)\big) = r$ and such that $r+1 \in \mathrm{supp}(x_n)$ if and only if $r+1 \in \mathrm{supp}(x_m)$. It is readily verified that $\min\big(\mathrm{supp}(x_n + x_m)\big) = r+1$. Hence $x_n + x_m \in A_{1-i}$. □

Theorem 2.12. *The collection of all infinite FS-big sets is not partition regular.*

Proof. In Theorem 3.1 in [14] it is shown that $\mathbb{T}\big|_1$ is infinite FS-big. The proof is based on the definition of the Thue–Morse word via binary expansions. The idea is similar to the proof of the fact that $\mathbb{T}\big|_0$ is an IP-set we saw in Example 1, but the proof of the fact that $\mathbb{T}\big|_1$ is infinite FS-big is more technical. The result now follows from Lemma 2.11. □

It follows from Lemma 2.11 and Theorem 2.7 that $\mathbb{T}\big|_1$ is not an IP-set. In fact, if $\mathbb{T}\big|_1$ were an IP-set, then by Theorem 2.7 one of the sets A_0 or A_1 in the proof of Lemma 2.11 would be an IP-set. But this contradicts Lemma 2.11. Similarly, since $\mathbb{T}\big|_1 \in \Sigma^\infty \subset \Sigma$ and Σ is partition regular, it follows that one of the two sets A_0, A_1 in the proof of Lemma 2.11 is in Σ. Thus we deduce that

$$\mathcal{IP} \subsetneq \Sigma^\infty \subsetneq \Sigma.$$

It follows from Theorems 2.7 and 2.8 that $\Sigma^* \subset \Sigma$ and $\mathcal{IP}^* \subset \mathcal{IP}$.

Theorem 2.7 can be used to show that a given set is not an IP-set. For instance we may now verify that $\mathbf{f}\big|_1$ is not an IP-set, and hence $\mathbf{f}\big|_0$ is an IP*-set. Let $\alpha = \frac{3-\sqrt{5}}{2}$. It is known that the Fibonacci word $\mathbf{f}$ is the orbit of the point α under

irrational rotation R_α on the unit circle by α. Let I be the interval $[1-\alpha, 1)$ (the interval coded by 1). So $n \in \mathbf{f}\big|_1$ if and only if $R_\alpha^n(\alpha) = \{\alpha + n\alpha\} = \{(n+1)\alpha\} \in I$. Fix

$$(1-\alpha)/3 \leq \alpha' \leq (1-\alpha)/2$$

and put

$$I_1 = [1-\alpha, 1-\alpha') \text{ and } I_2 = [1-\alpha', 1).$$

Since $\alpha' \leq (1-\alpha)/2$ it follows that $\alpha' < \alpha$. Also for $j = 1, 2$ set

$$A_j = \{n \in \mathbb{N} \mid R^n(\alpha) \in I_j\}.$$

Thus A_1, A_2 partitions the set $\mathbf{f}\big|_1$. Using Theorem 2.7 we now show that $\mathbf{f}\big|_1$ is not an IP-set by showing that the sum of any three elements of A_1 belongs to $\mathbf{f}\big|_0$ and that the sum of any two elements of A_2 belongs to $\mathbf{f}\big|_0$.

Now take any $n_1, n_2, n_3 \in A_1$ and set

$$x_1 = \{(n_1+1)\alpha\}, \;\; x_2 = \{(n_2+1)\alpha\}, \;\; x_3 = \{(n_3+1)\alpha\}.$$

Then $x_1, x_2, x_3 \in [1-\alpha, 1-\alpha')$ and $n_1 + n_2 + n_3$ corresponds to the point

$$\{(n_1+n_2+n_3+1)\alpha\} = \{x_1 + x_2 + x_3 - 2\alpha\}.$$

Since $x_1, x_2, x_3 \in [1-\alpha, 1-\alpha')$, we have

$$\{x_1 + x_2 + x_3 - 2\alpha\} \in [\{3-5\alpha\}, \{3-3\alpha'-2\alpha\}).$$

Since $\alpha' \geq \frac{1-\alpha}{3}$ it follows that

$$2 - 3\alpha' - 2\alpha \leq 1 - \alpha,$$

and hence

$$\{2 - 3\alpha' - 2\alpha\} \leq 1 - \alpha,$$

which gives

$$\{3 - 3\alpha' - 2\alpha\} \leq 1 - \alpha$$

as required.

Similarly take any $n_1, n_2 \in A_2$. Set

$$x_1 = \{(n_1+1)\alpha\}, \;\; x_2 = \{(n_2+1)\alpha\}$$

so that $x_1, x_2 \in [1-\alpha', 1)$. Then $n_1 + n_2$ corresponds to the point

$$\{(n_1+n_2+1)\alpha\} = \{x_1 + x_2 - \alpha\}.$$

Since $x_1, x_2 \in [1-\alpha', 1)$, we have

$$\{x_1 + x_2 - \alpha\} \in [\{2 - 2\alpha' - \alpha\}, 1-\alpha).$$

Since

$$\alpha' \leq \frac{1-\alpha}{2}$$

it follows that

$$\{1 - 2\alpha' - \alpha\} \geq 0,$$

and hence

$$\{2 - 2\alpha' - \alpha\} \geq 0.$$

2.3. Central sets

We consider a variation of Proposition 2.4. Let $\mathbb{A}$ be a finite non-empty set and let $T : \mathbb{A}^{\mathbb{N}} \to \mathbb{A}^{\mathbb{N}}$ denote the shift map.Two infinite words ω and ν are said to be *proximal* if for every positive integer n there exists a nonnegative integer m such that $T^m(\omega)$ and $T^m(\nu)$ share a common prefix of length n. In other words, ω and ν have arbitrarily long sections of agreement or coincidence.

Proposition 2.13. *Let $\omega, \nu \in \mathbb{A}^{\mathbb{N}}$. Suppose that ν is uniformly recurrent and ω is proximal to ν. Let $a = \nu_0$ denote the first letter of ν. Then $\omega\big|_a$ is an IP-set.*

Proof. As in the proof of Proposition 2.4, we recursively construct an increasing sequence $0 = x_0 < x_1 < x_2 < \cdots$ and prefixes $(B_k)_{k=0}^{\infty}$ of ν with $B_0 = \nu_0 = a$, $|B_k| = x_{k-1} + x_{k-2} + \cdots + x_1 + 1$ and where x_k is an occurrence of B_{k-1}. However, for each $k \geq 1$, we pick x_k in such a way that the occurrence B_{k-1} at x_k is completely contained in a region of agreement between ν and ω. This is possible since ω and ν agree on arbitrarily long stretches and ν is uniformly recurrent. Whence for each $k \geq 1$, the prefix B_{k-1} of ν occurs also in ω in position x_k. Hence $FS(x_t)_{t=1}^{\infty} \subset \omega\big|_a$. □

Given a set $A \subset \mathbb{N}$, we denote by $\chi(A) \in \{0,1\}^{\mathbb{N}}$ the word whose nth coordinate is equal to 1 if $n \in A$ and 0 otherwise. In other words, $A = \chi(A)\big|_1$.

Definition 2.14. A subset A of $\mathbb{N}$ is called a *central* set if $\chi(A)$ is proximal to a uniformly recurrent word beginning in 1. We say A is *central** if $A \cap B \neq \emptyset$ for every central set $B \subseteq \mathbb{N}$.

Let $\mathcal{C}$ denote the set of all central sets. Then it follows from Proposition 2.13 that $\mathcal{C} \subset \mathcal{IP}$. We recall that for the Sierpinski word $\mathbf{s}$ the set $\mathbf{s}\big|_0$ is an IP-set. On the other hand, it is readily verified that $\mathbf{s}$ is not proximal to any uniformly recurrent word beginning in 0 so that $\mathbf{s}\big|_0$ is not a central set. Whence $\mathcal{C} \subsetneq \mathcal{IP}$. On the other hand $\mathbf{s}$ is proximal to the uniformly recurrent word 1^{∞}, whence $\mathbf{s}\big|_1$ is a central set and hence an IP-set.
Central sets were originally defined by Furstenberg in [24] using notions from topological dynamics:

Definition 2.15. A subset $A \subset \mathbb{N}$ is called *central* if there exists a compact metric space (X, d) and a continuous map $T : X \to X$, points $x, y \in X$ and a neighborhood U of y such that

- y is a uniformly recurrent point in X,
- x and y are proximal,
- $A = \{n \in \mathbb{N} \,|\, T^n(x) \in U\}$.

Recall that in this more general setting, $x \in X$ is said to be *uniformly recurrent* in X if for every neighborhood V of x the set $\{n \,|\, T^n(x) \in V\}$ is syndetic. Also two points $x, y \in X$ are said to be *proximal* if for every $\epsilon > 0$ there exists $n \in \mathbb{N}$ such that $d(T^n(x), T^n(y)) < \epsilon$. It is evident that if $A \subset \mathbb{N}$ is central in the sense of

Definition 2.14 then it is central in the sense of Definition 2.15. It could be verified that in fact the two definitions are equivalent.

Central sets are known to have substantial combinatorial structure. For example, any central set contains arbitrarily long arithmetic progressions, and solutions to all partition regular systems of homogeneous linear equations (see for example [10]). Many of the rich properties of central sets are a consequence of the *Central Sets Theorem* first proved by Furstenberg in Proposition 8.21 in [24] (see also [16, 10, 28]). Furstenberg pointed out that as an immediate consequence of the Central Sets Theorem one has that whenever $\mathbb{N}$ is divided into finitely many classes, and a sequence $(x_n)_{n\in\mathbb{N}}$ is given, one of the classes must contain arbitrarily long arithmetic progressions whose increment belongs to $\{\sum_{n\in F} x_n | F \in \mathrm{Fin}\}$. We shall see in the next section that $\mathcal{C}$ is also partition regular. We shall also see that $\mathcal{C}$ is closed under supersets (i.e., any set containing central set is central), which is not obvious from the above definitions.

3. Additive Properties through the Stone–Čech compactification of $\mathbb{N}$

3.1. Stone–Čech compactification

In this section we re-visit IP-sets and central sets via the algebraic/topological properties of the Stone–Čech compactification of $\mathbb{N}$, denoted $\beta\mathbb{N}$. We regard $\beta\mathbb{N}$ as the set of all ultrafilters on $\mathbb{N}$ with the *Stone topology.*

Definition 3.1. A collection $\mathcal{U}$ of subsets of $\mathbb{N}$ is called an *ultrafilter* if the following conditions hold:

- $\emptyset \notin \mathcal{U}$.
- If $A \in \mathcal{U}$ and $A \subseteq B$, then $B \in \mathcal{U}$.
- $A \cap B \in \mathcal{U}$ whenever both A and B belong to $\mathcal{U}$.
- For every $A \subseteq \mathbb{N}$ either $A \in \mathcal{U}$ or $A^c \in \mathcal{U}$ where A^c denotes the complement of A.

For each natural number $n \in \mathbb{N}$, it is readily verified that the set $\mathcal{U}_n = \{A \subseteq \mathbb{N} \,|\, n \in A\}$ is an ultrafilter. This defines an injection $i : \mathbb{N} \hookrightarrow \beta\mathbb{N}$ by: $n \mapsto \mathcal{U}_n$. An ultrafilter of this form is said to be *principal.* By way of Zorn's lemma, one can show the existence of non-principal (or *free*) ultrafilters.

It is customary to denote elements of $\beta\mathbb{N}$ by letters $p, q, r \ldots$. For each set $A \subseteq \mathbb{N}$, we set $A^\circ = \{p \in \beta\mathbb{N} | A \in p\}$. Then the set $\mathcal{B} = \{A^\circ | A \subseteq \mathbb{N}\}$ forms a basis for the open sets (as well as a basis for the closed sets) of $\beta\mathbb{N}$ and defines a topology on $\beta\mathbb{N}$ with respect to which $\beta\mathbb{N}$ is both compact and Hausdorff. Although the existence of free ultrafilters requires Zorn's lemma, the cardinality of $\beta\mathbb{N}$ is $2^{2^{\mathbb{N}}}$ from which it follows that $\beta\mathbb{N}$ is not metrizable.

There is a natural extension of the operation of addition $+$ on $\mathbb{N}$ to $\beta\mathbb{N}$ making $\beta\mathbb{N}$ a compact *left-topological semigroup.* More precisely we define addition of two

ultrafilters p, q by the following rule:

$$p + q = \{A \subseteq \mathbb{N} \mid \{n \in \mathbb{N} | A - n \in p\} \in q\}.$$

It is readily verified that $p + q$ is once again an ultrafilter and that for each fixed $p \in \beta\mathbb{N}$, the mapping $q \mapsto p + q$ defines a continuous map from $\beta\mathbb{N}$ into itself.[1] The operation of addition in $\beta\mathbb{N}$ is associative and for principal ultrafilters we have $\mathcal{U}_m + \mathcal{U}_n = \mathcal{U}_{m+n}$. However in general addition of ultrafilters is highly non-commutative. In fact it can be shown that the center is precisely the set of all principal ultrafilters [27].

3.2. Idempotent ultrafilters

Let $(\mathcal{S}, +)$ be a semigroup. An element $p \in \mathcal{S}$ is called an *idempotent* if $p + p = p$. We recall the following result of Ellis [23]:

Theorem 3.2 (Ellis [23]). *Let $(\mathcal{S}, +)$ be a compact left-topological semigroup (i.e., $\forall x \in \mathcal{S}$ the mapping $y \mapsto x + y$ is continuous). Then $\mathcal{S}$ contains an idempotent.*

It follows that $\beta\mathbb{N}$ contains a non-principal ultrafilter p satisfying $p + p = p$. In fact, we could simply apply Ellis's result to the semigroup $\beta\mathbb{N} - \mathcal{U}_0$. This would then exclude the only principal idempotent ultrafilter, namely $\mathcal{U}_0$. From here on, by an idempotent ultrafilter in $\beta\mathbb{N}$ we mean a free idempotent ultrafilter.

Hindman proved the following striking result linking IP-sets and idempotents in $\beta\mathbb{N}$:

Theorem 3.3 ([27, Thm. 5.12]). *A subset $A \subseteq \mathbb{N}$ is an IP-set if and only if $A \in p$ for some idempotent $p \in \beta\mathbb{N}$.*

It follows immediately that A is an IP*-set if and only if $A \in p$ for every idempotent $p \in \beta\mathbb{N}$ (see Theorem 2.15 in [8]). We also note that it follows from Theorem 3.3 that the collection of all IP-sets is partition regular.

We illustrate one direction of Theorem 3.3. Suppose $p \in \beta\mathbb{N}$ is an idempotent ultrafilter and $A \in p$. We will show that A is an IP-set. Put $A_0 = A$. Then we have

$$\begin{aligned} A_0 \in p &\Leftrightarrow A_0 \in p + p \\ &\Leftrightarrow \{n \mid A_0 - n \in p\} \in p \\ &\Leftrightarrow \exists x_0 \in A_0 \text{ with } A_0 - x_0 \in p \\ &\Leftrightarrow \exists x_0 \in A_0 \text{ with } (A_0 - x_0) \cap A_0 \in p. \end{aligned}$$

Set $A_1 = (A_0 - x_0) \cap A_0$. Since $A_1 \in p$, as above we can find $x_1 \in A_1$ with $x_1 > x_0$ and $(A_1 - x_1) \cap A_1 \in p$. Set $A_2 = (A_1 - x_1) \cap A_1$. Again since $A_2 \in p$ there exists $x_2 \in A_2$ with $x_2 > x_1$ and $(A_2 - x_2) \cap A_2 \in p$. Continuing in this way we produce an increasing sequence $x_0 < x_1 < x_2 < \cdots$ and a nested sequence $A_0 \supset A_1 \supset A_2 \supset \cdots$ belonging to p and such that $x_k \in A_k$ and $A_{k+1} = (A_k - x_k) \cap A_k$ for each $k \geq 0$ (see Figure 3).

[1] Our definition of addition of ultrafilters is the same as that given in [8] but is the reverse of that given in [27] in which $A \in p + q$ if and only if $\{n \in \mathbb{N} | A - n \in q\} \in p\}$. In this case, $\beta\mathbb{N}$ becomes a compact right-topological semigroup.

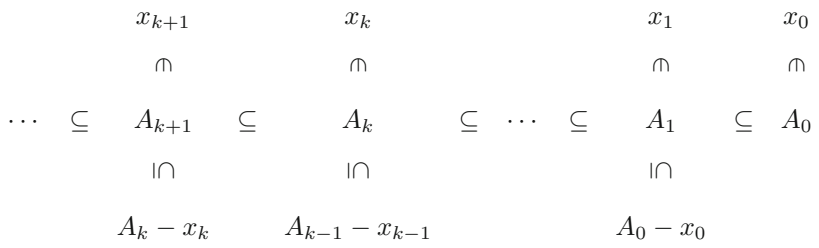

FIGURE 3. Proof of Theorem 3.3.

It is now easy to check that $FS(x_t)_{t=0}^{\infty} \subset A$. For instance, let us verify that $x_{11} + x_5 + x_2 + x_1 \in A$.

$$x_{11} \in A_{11} \subset A_6 \subset A_5 - x_5$$

whence

$$x_{11} + x_5 \in A_5 \subset A_3 \subset A_2 - x_2$$

whence

$$x_{11} + x_5 + x_2 \in A_2 \subset A_1 - x_1$$

whence

$$x_{11} + x_5 + x_2 + x_1 \in A_1 \subset A_0 = A.$$

Central sets may also be alternatively defined in terms of belonging to a special class of free ultrafilters, called minimal idempotents. Let $(\mathcal{S}, +)$ be any semigroup. Recall that a subset $\mathcal{I} \subseteq \mathcal{S}$ is called a *right (resp. left) ideal* if $\mathcal{I}+\mathcal{S} \subseteq \mathcal{I}$ (resp. $\mathcal{S}+\mathcal{I} \subseteq \mathcal{I}$). It is called a *two-sided ideal* if it is both a left and right ideal. A right (resp. left) ideal $\mathcal{I}$ is called *minimal* if every right (resp. left) ideal $\mathcal{J}$ included in $\mathcal{I}$ coincides with $\mathcal{I}$.

Minimal right/left ideals do not necessarily exist, e.g., the commutative semigroup $(\mathbb{N}, +)$ has no minimal right/left ideals (the ideals in $\mathbb{N}$ are all of the form $\mathcal{I}_n = [n, +\infty) = \{m \in \mathbb{N} \,|\, m \geq n\}$.) However, every compact Hausdorff left-topological semigroup $\mathcal{S}$ (e.g., $\beta\mathbb{N}$) admits a smallest two-sided ideal $K(\mathcal{S})$ which is at the same time the union of all minimal right ideals of $\mathcal{S}$ and the union of all minimal left ideals of $\mathcal{S}$ (see for instance [27]). It is readily verified that the intersection of any minimal left ideal with any minimal right ideal is a group. In particular, there are idempotents in $K(\mathcal{S})$. Such idempotents are called minimal and their elements are called central sets:

Definition 3.4. An idempotent p is called a *minimal* idempotent of $\mathcal{S}$ if it belongs to $K(\mathcal{S})$.

Definition 3.5. A subset $A \subset \mathbb{N}$ is called *central* if it is a member of some minimal idempotent in $\beta\mathbb{N}$. It is called a central*-set if it belongs to every minimal idempotent in $\beta\mathbb{N}$.

The equivalence between definitions 3.5 and 2.15 is due to Bergelson and Hindman in [9]. It follows from the above definition that every central set is an IP-set and that the property of being central is partition regular. This definition also implies that central sets are closed under supersets.

3.3. Limits of ultrafilters

It is often convenient to think of an ultrafilter p as a $\{0,1\}$-valued, finitely additive probability measure on the power set of $\mathbb{N}$. More precisely, for any subset $A \subseteq \mathbb{N}$, we say A has p-measure 1, or is p-large if $A \in p$. This notion of measure gives rise to a notion of convergence of sequences indexed by $\mathbb{N}$ which is the key tool in allowing us to apply ideas from combinatorics on words to the framework of ultrafilters. However, from our point of view, it is more natural to define it alternatively as a mapping from words to words (see Remark 3.7). Let $\mathbb{A}$ denote a non-empty finite set. Then each ultrafilter $p \in \beta\mathbb{N}$ naturally defines a mapping

$$p^* : \mathbb{A}^{\mathbb{N}} \to \mathbb{A}^{\mathbb{N}}$$

as follows:

Definition 3.6. For each $p \in \beta\mathbb{N}$ and $\omega \in \mathbb{A}^{\mathbb{N}}$, we define $p^*(\omega) \in \mathbb{A}^{\mathbb{N}}$ by the condition: $u \in \mathbb{A}^*$ is a prefix of $p^*(\omega) \Leftrightarrow \omega\big|_u \in p$.

We note that if $u, v \in \mathbb{A}^*$, $\omega\big|_u, \omega\big|_v \in p$ and $|v| \geq |u|$, then u is a prefix of v. In fact, if v' denotes the prefix of v of length $|u|$ then as $\omega\big|_v \subseteq \omega\big|_{v'}$, it follows that $\omega\big|_{v'} \in p$ and hence $u = v'$. Thus $p^*(\omega)$ is well defined.

Remark 3.7. It is readily verified that our definition of p^* coincides with that of $p\text{-lim}_n$. More precisely, given a sequence $(x_n)_{n\in\mathbb{N}}$ in a topological space and an ultrafilter $p \in \beta\mathbb{N}$, we write $p\text{-lim}_n x_n = y$ if for every neighborhood U_y of y one has $\{n \,|\, x_n \in U_y\} \in p$. In our case we have $p^*(\omega) = p\text{-lim}_n(T^n(\omega))$ (see [27]).

We note that if $\omega, \nu \in \mathbb{A}^{\mathbb{N}}$ and if each prefix u of ν is a factor of ω, then there exists an ultrafilter $p \in \beta\mathbb{N}$ such that $p^*(\omega) = \nu$. In fact, the set

$$\mathcal{D} = \left\{\omega\big|_u \,|\, u \text{ is a prefix of } \nu\right\}$$

satisfies the finite intersection property, and hence by a routine argument involving Zorn's lemma it follows that there exists a $p \in \beta\mathbb{N}$ with $\mathcal{D} \subseteq p$.

It follows immediately from the definition of p^*, Definition 3.5 and Theorem 3.3 that

Lemma 3.8. *The set $\omega\big|_u$ is an IP-set (resp. central set) if and only if u is a prefix of $p^*(\omega)$ for some idempotent (resp. minimal idempotent) $p \in \beta\mathbb{N}$.*

Lemma 3.9. *For each $p \in \beta\mathbb{N}$, $\omega \in \mathbb{A}^{\mathbb{N}}$ and $u \in \mathbb{A}^*$ we have*

$$p^*(\omega)\big|_u = \left\{m \in \mathbb{N} \,|\, \omega\big|_u - m \in p\right\}$$

where $\omega\big|_u - m$ is defined as the set of all $n \in \mathbb{N}$ such that $n + m \in \omega\big|_u$.

Proof. Suppose $m \in p^*(\omega)\big|_u$. Then by definition u occurs in position m in $p^*(\omega)$. Let v denote the prefix of $p^*(\omega)$ of length $|v| = m + |u|$. Then, as u is a suffix of v we have $\omega\big|_v + m \subseteq \omega\big|_u$ and hence $\omega\big|_v \subseteq \omega\big|_u - m$. But as v is a prefix of $p^*(\omega)$ we have $\omega\big|_v \in p$ and hence $\omega\big|_u - m \in p$ as required.

Conversely, fix $m \in \mathbb{N}$ such that $\omega\big|_u - m \in p$. Let Z be the set of all factors v of ω of length $|v| = m + |u|$ ending in u. Then

$$\omega\big|_u - m \subseteq \bigcup_{v \in Z} \omega\big|_v.$$

It follows that there exists $v \in Z$ such that $\omega\big|_v \in p$. In other words, there exists $v \in Z$ such that v is a prefix of $p^*(\omega)$. It follows that u occurs in position m in $p^*(\omega)$. □

Lemma 3.10. *For $p, q \in \beta\mathbb{N}$ and $\omega \in \mathbb{A}^{\mathbb{N}}$, we have $(p+q)^*(\omega) = q^*(p^*(\omega))$. In particular, if p is an idempotent, then $p^*(p^*(\omega)) = p^*(\omega)$.*

Proof. For each word $u \in \mathbb{A}^*$ we have that u is a prefix of $(p+q)^*(\omega)$ if and only if

$$\omega\big|_u \in p + q \Leftrightarrow \{m \in \mathbb{N} \,|\, \omega\big|_u - m \in p\} \in q.$$

On the other hand, u is a prefix of $q^*(p^*(\omega))$ if and only if $p^*(\omega)\big|_u \in q$. The result now follows immediately from the preceding lemma. □

Lemma 3.11. *For each $p \in \beta\mathbb{N}$ and $\omega \in \mathbb{A}^{\mathbb{N}}$ we have $p^*(T(\omega)) = T(p^*(\omega))$ where $T : \mathbb{A}^{\mathbb{N}} \to \mathbb{A}^{\mathbb{N}}$ denotes the shift map.*

Proof. Assume $u \in \mathbb{A}^*$ is a prefix of $p^*(T(\omega))$. Then $T(\omega)\big|_u \in p$. But

$$T(\omega)\big|_u = \bigcup_{a \in \mathbb{A}} \omega\big|_{au}.$$

It follows that there exists $a \in \mathbb{A}$ such that $\omega\big|_{au} \in p$. Thus au is a prefix of $p^*(\omega)$ and hence u is a prefix of $T(p^*(\omega))$. □

We will make use of the following key result in [27] (see also [13, Thm. 1] and [8, Thm. 3.4]):

Theorem 3.12 ([27, Thm. 19.26]). *Given two infinite words $x, y \in \mathbb{A}^{\infty}$, there is a minimal idempotent $p \in \beta\mathbb{N}$ such that $p^*(x) = y$ if and only if x and y are proximal and y is uniformly recurrent.*

As a consequence we have

Theorem 3.13. *Let $\omega \in \mathbb{A}^{\mathbb{N}}$ be a uniformly recurrent word, and let $u \in \mathbb{A}^+$. Then $\omega\big|_u$ is an IP-set if and only if $\omega\big|_u$ is a central set.*

Proof. For any $A \subset \mathbb{N}$ we have that if A is central then A belongs to some minimal idempotent $p \in \beta\mathbb{N}$ and hence in particular A belongs to an idempotent in $\beta\mathbb{N}$. Hence by Theorem 3.3 we have that A is an IP-set. Now suppose that $\omega\big|_u$ is an IP-set. Then $\omega\big|_u$ belongs to some idempotent $p \in \beta\mathbb{N}$. Set $\nu = p^*(\omega)$. Then u is a prefix of ν. Also, since p is idempotent we have $p^*(\nu) = p^*(p^*(\omega)) = p^*(\omega) = \nu$. Hence

for every prefix v of ν we have that $\nu\big|_v \in p$ and $\omega\big|_v \in p$ and hence $\nu\big|_v \cap \omega\big|_v \in p$. In particular $\nu\big|_v \cap \omega\big|_v \neq \emptyset$. Hence ω and ν are proximal. Since ω is uniformly recurrent, it follows that ν is also uniformly recurrent. Hence by Theorem 3.12 there exists a minimal idempotent q with $q^*(\omega) = \nu$. Hence $\omega\big|_u \in q$, whence $\omega\big|_u$ is central. □

This allows us to simultaneously state our results in terms of IP-sets and central sets.

4. Central sets defined by some families of uniformly recurrent words

In this section we will consider several families of uniformly recurrent words and study combinatorial properties of central sets defined by words from these families.

First we give a brief summary of some of the basic background on subshifts. Given an infinite word $\omega \in \mathbb{A}^{\mathbb{N}}$, we denote by $\mathcal{F}_\omega(n)$ the set of all factors of ω of length n, and set

$$\mathcal{F}_\omega = \bigcup_{n\in\mathbb{N}} \mathcal{F}_\omega(n).$$

We endow $\mathbb{A}^{\mathbb{N}}$ with the topology generated by the metric

$$d(x,y) = \frac{1}{2^n} \text{ where } n = \inf\{k : x_k \neq y_k\}$$

whenever $x = (x_n)_{n\in\mathbb{N}}$ and $y = (y_n)_{n\in\mathbb{N}}$ are two elements of $\mathbb{A}^{\mathbb{N}}$. By a *subshift* on $\mathbb{A}$ we mean a pair (X,T) where X is a closed and T-invariant subset of $\mathbb{A}^{\mathbb{N}}$, and T is the shift transformation map. A subshift (X,T) is said to be *minimal* whenever X and the empty set are the only T-invariant closed subsets of X. To each $\omega \in \mathbb{A}^{\mathbb{N}}$ is associated the subshift (X,T) where X is the shift orbit closure of ω. If ω is uniformly recurrent, then the associated subshift (X,T) is minimal. Thus any two words x and y in X have exactly the same set of factors, i.e., $\mathcal{F}_x = \mathcal{F}_y$. In this case we denote by $\mathcal{F}_X$ the set of factors of any word $x \in X$.

4.1. Sturmian partitions & central sets

In this subsection we study the additive properties of sets defined by the family of Sturmian words. More precisely, we obtain a complete characterization of which factors of a Sturmian word give rise to central sets.

Let $\omega \in \mathbb{A}^{\mathbb{N}}$ and set

$$\rho_\omega(n) = \text{Card}(\mathcal{F}_\omega(n)).$$

The function $\rho_\omega : \mathbb{N} \to \mathbb{N}$ is called the *factor complexity function* of ω. Given a minimal subshift (X,T) on A, we have $\mathcal{F}_\omega(n) = \mathcal{F}_{\omega'}(n)$ for all $\omega, \omega' \in X$ and $n \in \mathbb{N}$. Thus we can define the factor complexity $\rho_{(X,T)}(n)$ of a minimal subshift (X,T) by

$$\rho_{(X,T)}(n) = \rho_\omega(n)$$

for any $\omega \in X$.

A word $\omega \in \mathbb{A}^{\mathbb{N}}$ is *periodic* if there exists a positive integer p such that $\omega_{i+p} = \omega_i$ for all indices i, and it is *ultimately periodic* if $\omega_{i+p} = \omega_i$ for all sufficiently large i. An infinite word is *aperiodic* if it is not ultimately periodic. By a celebrated result due to Hedlund and Morse [32], a word is ultimately periodic if and only if its factor complexity is uniformly bounded. In particular, $p_\omega(n) < n$ for all n sufficiently large. Words whose factor complexity $\rho_\omega(n) = n + 1$ for all $n \geq 0$ are called *Sturmian words*. Thus, Sturmian words are those aperiodic words having the lowest complexity. Since $\rho_\omega(1) = 2$, it follows that Sturmian words are binary words. The most extensively studied Sturmian word is the Fibonacci word $\mathbf{f}$ defined in Subsection 2.1.

A factor u of an infinite word $\omega \in \mathbb{A}^{\mathbb{N}}$ is called *right special* if both ua and ub are factors of ω for some pair of distinct letters $a, b \in \mathbb{A}$. Similarly u is called *left special* if both au and bu are factors of ω for some pair of distinct letters $a, b \in \mathbb{A}$. The factor u is called *bispecial* if it is both right special and left special. Let $\omega \in \{0,1\}^{\mathbb{N}}$ be a Sturmian word, and let Ω denote the shift orbit closure of ω. The condition $\rho_\omega(n) = n+1$ implies the existence of exactly one right special and one left special factor of each length. Clearly, given any two left special factors, one is necessarily a prefix of the other. It follows that Ω contains a unique word all of whose prefixes are left special factors of ω. Such a word is called the *characteristic word* and denoted $\tilde{\omega}$. It follows that both $0\tilde{\omega}, 1\tilde{\omega} \in \Omega$. It is readily verified that the Fibonacci word above is a characteristic Sturmian word. A Sturmian word ω is called *singular* if $T^n(\omega) = \tilde{\omega}$ for some $n \geq 1$. Otherwise it is said to be *nonsingular.*

The next two theorems give a complete characterization of those factors u of a Sturmian word $\omega \in \{0,1\}^{\mathbb{N}}$ for which $\omega\big|_u$ is an IP-set (respectively central set).

Theorem 4.1 ([15]). *Let $\omega \in \Omega$ be a nonsingular Sturmian word, and u a factor of ω. Then $\omega\big|_u$ is an IP-set (resp. central set) if and only if u is a prefix of ω.*

As a consequence we obtain:

Corollary 4.2. *For every prefix v of a nonsingular Sturmian word ω and $n \in \omega\big|_v$, the set $\omega\big|_v - n$ is an IP*-set (resp. central* set).*

We note that in general the property of being an IP*-set is not translation invariant. See also Theorem 1.1 in [10].

Theorem 4.3 ([15]). *Let $\omega \in \Omega$ be a Sturmian word such that $T^{n_0}(\omega) = \tilde{\omega}$ with $n_0 \geq 1$. Then $\omega\big|_u$ is an IP-set (or central set) if and only if either u is a prefix of ω or a prefix of ω' where ω' is the unique other element of Ω with $T^{n_0}(\omega') = \tilde{\omega}$.*

To prove the theorems, we need several lemmas.

Lemma 4.4. *If $\omega, \omega', \omega'' \in \Omega$ are such that $T^{n_0}(\omega) = T^{n_0}(\omega') = T^{n_0}(\omega'')$, then Card$\{\omega, \omega', \omega''\} \leq 2$.*

Proof. This follows immediately from the fact that Ω contains a unique characteristic word and that this word is aperiodic. □

We will make use of the following key lemma which essentially says that two distinct Sturmian words ω and ω' are proximal if and only if $T^n(\omega) = T^n(\omega') = \tilde{\omega}$ for some $n \geq 1$.

Lemma 4.5. *Let ω and ω' be distinct elements of Ω. Then either $T^n(\omega) = T^n(\omega') = \tilde{\omega}$ for some $n \geq 1$, or there exists $N > 0$ such that*

$$\omega_n\omega_{n+1}\cdots\omega_{n+N} \neq \omega'_n\omega'_{n+1}\cdots\omega'_{n+N}$$

for every $n \in \mathbb{N}$.

The proof of this lemma is based on the definition of Sturmian words via rotations. For the details of the proof we refer to [15].
We first consider the case of nonsingular Sturmian words:

Lemma 4.6. *Let $\omega \in \{0,1\}^{\mathbb{N}}$ be a nonsingular Sturmian word and $p \in \beta\mathbb{N}$ an idempotent ultrafilter. Then $p^*(\omega) = \omega$.*

Proof. Suppose to the contrary that $p^*(\omega) \neq \omega$. Then since ω is nonsingular, Lemma 4.5 implies that for all sufficiently long factors u of ω, we have that $\omega\big|_u \cap p^*(\omega)\big|_u = \emptyset$. But, by Lemma 3.10 we have $p^*(p^*(\omega)) = p^*(\omega)$, that is the image under p^* of ω and $p^*(\omega)$ coincides. It follows by definition of p^* that for every prefix u of $p^*(\omega)$ we have $\omega\big|_u \in p$ and $p^*(\omega)\big|_u \in p$ and hence $\omega\big|_u \cap p^*(\omega)\big|_u \in p$, a contradiction. □

Proof of Theorem 4.1. Let ω be a nonsingular Sturmian word, u a prefix of ω, and $p \in \beta\mathbb{N}$ an idempotent ultrafilter. Then by Lemma 4.6 u is a prefix of $p^*(\omega)$ and hence $\omega\big|_u \in p$. Thus for each prefix u of ω the set $\omega\big|_u$ belongs to every idempotent ultrafilter and hence is an IP*-set. It follows that if $v \in F$ is not a prefix of ω, then $\omega\big|_v$ is not an IP-set. Finally, let v be any factor of ω and $n \in \mathbb{N}$. Then $\omega\big|_v - n = T^n(\omega)\big|_v$. If $n \in \omega\big|_v$, then v is a prefix of $T^n(\omega)$ from which it follows that

$$\omega\big|_v - n = T^n(\omega)\big|_v \in p.$$

Hence $\omega\big|_v - n$ is an IP*-set □

As a consequence of the above theorem we have

Corollary 4.7. *Let ω and ω' be two nonsingular Sturmian words, not necessarily of the same slope. Then for every prefix u of ω and every prefix u' of ω' we have that $\omega\big|_u \cap \omega'\big|_{u'}$ is an IP*-set (resp. central* set), in particular the intersection is infinite.*

We note that the assumption that ω and ω' be nonsingular is necessary, as for example if we consider $\omega = 0\mathbf{f}$ and $\omega' = 1\mathbf{f}$ with $\mathbf{f}$ the Fibonacci word, then $\omega\big|_0 \cap \omega'\big|_1 = \{0\}$.

Proof. Let ω and ω' be two nonsingular Sturmian words, u a prefix of ω, u' a prefix of ω', and $p \in \beta\mathbb{N}$ an idempotent ultrafilter. Then by Theorem 4.1 we have that $\omega\big|_u \in p$ and $\omega\big|_{u'} \in p$ and hence $\omega\big|_u \cap \omega\big|_{u'} \in p$. Thus $\omega\big|_u \cap \omega\big|_{u'}$ belongs to every idempotent and hence is an IP*-set. □

We next consider singular Sturmian words.

Lemma 4.8. *Let $\omega,\omega' \in \Omega$ be distinct Sturmian words such that $T^{n_0}(\omega) = T^{n_0}(\omega') = \tilde{\omega}$ for some $n_0 \geq 1$. Then for every $u \in \mathcal{F}$ and every non-principal ultrafilter $p \in \beta\mathbb{N}$ we have*

$$\omega\big|_u \in p \iff \omega'\big|_u \in p.$$

In particular, $p^*(\omega) = p^*(\omega')$.

Proof. Since p is a non-principal ultrafilter, we have that $\omega\big|_u \in p \iff \omega\big|_u \cap [N, +\infty) \in p$ for all $N \geq 1$. Similarly $\omega'\big|_u \in p \iff \omega'\big|_u \cap [N, +\infty) \in p$ for all $N \geq 1$. But for each $u \in \mathcal{F}$, we have $\omega\big|_u \cap [n_0, +\infty) = \omega'\big|_u \cap [n_0, +\infty)$. The result now follows. □

Lemma 4.9. *Let $\omega, \omega' \in \Omega$ be as in the previous lemma, and let $p \in \beta\mathbb{N}$ be an idempotent ultrafilter. Then $p^*(\omega) = p^*(\omega') \in \{\omega, \omega'\}$.*

Proof. That $p^*(\omega) = p^*(\omega')$ follows from the previous lemma and the fact that idempotent ultrafilters are non-principal (see for instance [8]). By Lemma 3.11, p^* commutes with the shift map T, and hence

$$T^{n_0} p^*(\omega) = p^*(T^{n_0}\omega) = p^*(\tilde{\omega}) = \tilde{\omega}$$

where the last equality follows from Lemma 4.6. By Lemma 4.4 applied to $\omega'' = p^*(\omega)$ it follows that $p^*(\omega) = \omega$ or $p^*(\omega) = \omega'$. □

Proof of Theorem 4.3. Let $\omega \in \Omega$ and n_0 be as in the statement of the theorem. Then there exists a unique $\omega' \in \Omega$ with $\omega' \neq \omega$ and with $T^{n_0}(\omega') = \tilde{\omega}$. Suppose that $\omega\big|_u$ is an IP-set for some $u \in \mathcal{F}$. Then by Lemma 3.8 it follows that u is a prefix of $p^*(\omega)$ for some idempotent ultrafilter $p \in \beta\mathbb{N}$. It follows from Lemma 4.9 that u is a prefix of ω or a prefix of ω'. This proves one direction.

To establish the other direction, we must show that $\omega\big|_u$ is a central set for each prefix u of ω or of ω'. By Theorem 3.12, there exist minimal idempotent ultrafilters $p_1, p_2 \in \beta\mathbb{N}$ such that $p_1^*(\omega) = \omega$ and $p_2^*(\omega) = \omega'$. The result now follows. □

4.2. Partitions defined by substitution rules

In this section we consider partitions defined by words generated by substitution rules. Let τ be a substitution on $\mathbb{A}$. A word $\omega \in \mathbb{A}^{\mathbb{N}}$ is called *periodic point* if $\tau^m(\omega) = \omega$ for some $m > 0$. A substitution τ is said to be *primitive* if there is a positive integer n such that for each pair $(i, j) \in \mathbb{A} \times \mathbb{A}$, the letter i occurs in $\tau^n(j)$. Although a primitive substitution τ may fail to have a fixed point, it has at least one periodic point. Associated to τ is the topological dynamical system (X, T),

where X is the shift orbit closure of a periodic point ω of τ. If τ is primitive, then (X, T) is independent of the choice of periodic point and is minimal.

By considering partitions of $\mathbb{N}$ defined by words generated by the generalized Thue–Morse substitution to an alphabet of size $r \geq 2$, we show the following:

Theorem 4.10 ([15]). *For each positive integer r let Ω denote the minimal subshift generated by the primitive constant length substitution $1 \mapsto 123\cdots r$, $2 \mapsto 23\cdots r1$, $\ldots$, $r \mapsto r12\cdots r-1$. Then for each positive integer N there exists $\omega \in \Omega$ such that the sets $A_i := \omega\big|_i$ for $1 \leq i \leq r$ define a partition of*

$$\mathbb{N} = A_1 \cup A_2 \cup \cdots \cup A_r$$

satisfying the following conditions:

- *$A_i - n$ is a central set for each $1 \leq i \leq r$ and $1 \leq n \leq N$.*
- *For each $n > N$, exactly one of the sets $\{A_1 - n, A_2 - n, \ldots, A_r - n\}$ is a central set.*

The second assertion of Theorem 4.10 relies on the fact that each fixed point of the generalized Thue–Morse substitution is distal. A point $x \in X$ is called *distal* if the only point in X proximal to x is x itself.

A minimal subshift (X, T) is said to be *topologically mixing* if for every any pair of factors $u, v \in \mathcal{F}_X$ there exists a positive integer N such that for each $n \geq N$, there exists a block of the form $uWv \in \mathcal{F}_X$ with $|W| = n$. A minimal subshift (X, T) is said to be *topologically weak mixing* if for every pair of factors $u, v \in \mathcal{F}_X$ the set

$$\{n \in \mathbb{N} \,|\, u\mathbb{A}^n v \cap \mathcal{F}_X \neq \emptyset\}$$

is thick, i.e., for every positive integer N, the set contains N consecutive positive integers.

By considering partitions defined by words generating minimal subshifts which are topologically weak mixing (for example the subshift generated by the substitution $0 \mapsto 001$ and $1 \mapsto 11001$) we prove that

Theorem 4.11 ([15]). *For each positive integer r there exists a partition of $\mathbb{N} = A_1 \cup A_2 \cup \cdots \cup A_r$ such that for each $1 \leq i \leq r$ and $n \geq 0$, the set $A_i - n$ is a central set.*

We remark that the proofs of Theorems 4.10 and 4.11 are both constructive, i.e., in each case we explicitly build a word defining the partition.

4.3. Infinite central partitions of $\mathbb{N}$

In this subsection we construct infinite partitions of $\mathbb{N}$ into central sets by using words on an infinite alphabet. Our construction makes use of the notion of *iterated palindromic closure operator* (first introduced in [17]):

Definition 4.12. The iterated palindromic operator ψ is defined inductively as follows:

- $\psi(\varepsilon) = \varepsilon$,
- For any word w and any letter a, $\psi(wa) = (\psi(w)a)^{(+)}$.

We denote with $w^{(+)}$ the *right palindromic closure* of the word w, i.e., the shortest palindrome which has w as a prefix.

For example, $\psi(aaba) = aabaaabaa$. The operator ψ has been extensively studied for its central role in constructing standard Sturmian and episturmian words. It follows immediately from the definition that if u is a prefix of v, then $\psi(u)$ is a prefix of $\psi(v)$. Thus, given an infinite word $\omega = \omega_0\omega_1\omega_2\cdots$ on the alphabet A we can define

$$\psi(\omega) = \lim_{n\to\infty} \psi(\omega_0\omega_1\omega_2\cdots\omega_n).$$

Proposition 4.13 ([15]). *Let $\omega = \psi(\Delta)$ where Δ is a right infinite word on an infinite alphabet $\mathbb{A}$ with the property that each letter $a \in \mathbb{A}$ occurs in Δ an infinite number of times. Then, for any $a \in \mathbb{A}$, the set $a\omega\big|_a$ is a central set, thus $\{\omega\big|_a + 1\}_{a\in\mathbb{A}}$ is an infinite partition of $\mathbb{N} - \{0\}$ into central sets*[2].

For the proof of this propositions we refer to [15].

5. Strong coincidence condition

Let $n \geq 2$ be a positive integer and set $\mathbb{A} = \{1, 2, \ldots, n\}$. Given a substitution $\tau : \mathbb{A} \to \mathbb{A}^+$, we consider the square matrix M_τ whose ijth entry is equal to $|\tau(j)|_i$, i.e., the number of occurrences of i in $\tau(j)$. We call this matrix the *Abelianization* of τ. Hence a substitution τ is primitive if all the entries of M_τ^n are strictly positive. In this case it is well known that the matrix M_τ has a simple positive Perron–Frobenius eigenvalue called the *dilation* of τ. A substitution τ is said to be *irreducible* if the minimal polynomial of its dilation is equal to the characteristic polynomial of its Abelianization M_τ. A substitution τ is said to be of *Pisot type* if its dilation is a Pisot number. Recall that a Pisot number is an algebraic integer greater than 1 all of whose algebraic conjugates lie strictly inside the unit circle.

A primitive substitution $\tau : \mathbb{A} \to \mathbb{A}^+$ is said to satisfy the *strong coincidence condition* if and only if any pair of fixed points x and y are *strongly coincident*, i.e., we can write $x = scx'$, and $y = tcy'$ for some $s, t \in \mathbb{A}^+$, $c \in \mathbb{A}$, and $x', y' \in \mathbb{A}^\infty$ with $s \sim_{\mathrm{ab}} t$. This combinatorial condition, originally due to P. Arnoux and S. Itô, is an extension of a similar condition considered by F.M. Dekking in [18] in the case of *constant length substitutions*, i.e., when $|\tau(a)| = |\tau(b)|$ for all $a, b \in \mathbb{A}$. Every such substitution τ has an algorithmically determined "pure base" substitution and Dekking proves that the strong coincidence condition is satisfied by the pure base if and only if the substitutive subshift associated with τ has *pure discrete spectrum*, i.e., is metrically isomorphic with translation on a compact Abelian group. The Thue–Morse substitution is equal to its pure base and clearly does not satisfy the strong coincidence condition – in fact the two fixed points disagree in each coordinate. It is conjectured however that if τ is an *irreducible* primitive

[2]This is a special case of a more general result of Hindman, Leader and Strauss [26] in which they show that every central set in $\mathbb{N}$ is a countable union of pairwise disjoint central sets.

substitution of Pisot type, then τ satisfies the strong coincidence condition. This conjecture is established for binary primitive substitutions of Pisot type in [3]. Otherwise the conjecture remains open for substitutions defined on alphabets of size greater than two. Substitutions of Pisot type provide a framework for non-constant length substitutions in which the strong coincidence condition is necessary (and, conjecturally, sufficient) for pure discrete spectrum (see [3, 4, 5, 6]). The following reformulation of the strong coincidence condition in terms of central sets was obtained by the second author together with M. Barge in [7]:

Theorem 5.1 ([7]). *Let τ be an irreducible primitive substitution of Pisot type. Then for any pair of fixed points x and y of τ the following are equivalent:*

1. *x and y are strongly coincident.*
2. *x and y are proximal.*
3. *There exists a minimal idempotent $p \in \beta\mathbb{N}$ such that $y = p^*(x)$.*
4. *For any prefix u of y, the set $x\big|_u$ is a central set.*

Remark 5.2. For a general primitive substitution we always have that (1) $\Longrightarrow$ (2) $\Longrightarrow$ (3) $\Longrightarrow$ (4). But in general in the non-Pisot setting, these conditions need not be equivalent: For instance, the two fixed points of the uniform substitution $a \mapsto aaab$, $b \mapsto bbab$ are proximal but do not satisfy the strong coincidence condition. V. Bergelson and Y. Son [12] showed that the fixed points of $a \mapsto aab$, $b \mapsto bbaab$ satisfy (4) but not (1), (2) and (3). It would be interesting to understand in general under what conditions do the idempotent ultrafilters permute the fixed points of substitutions.

Outline of the proof of Theorem 5.1. We first show that (1) $\Longrightarrow$ (2) $\Longrightarrow$ (3) $\Longrightarrow$ (4). Clearly (2) is immediate from the definition of strong coincidence. By Theorem 3.12 we have that (2) implies (3) and hence (4). The proof that (4) $\Longrightarrow$ (1) in [7] relies heavily on the machinery of "strand space" (a convenient presentation of tiling space) which will allow us to apply results developed for the $\mathbb{R}$-action on strand space to the shift action on words. Using strands and prior results of Barge and Kwapisz, the authors prove the following key proposition

Proposition 5.3 ([7, Prop. 4.8]). *Let $x, y \in X$. Then:*

1. *$\{x' \in X \mid x'$ is proximal with $x\}$ is finite; and*
2. *if x and y are proximal and fixed by τ, then x and y are strongly coincident.*

Proof. To see that (4) $\Longrightarrow$ (1), let x and y be fixed points of τ and suppose that for every prefix u of y the set $x\big|_u$ is a central set. This means that for every prefix u of y the set $x\big|_u$ belongs to some minimal idempotent $p_u \in \beta\mathbb{N}$. The collection

$$\mathcal{P} = \{p_u^*(x) \mid u \text{ is a prefix of } y\}$$

consists of infinite words in X each proximal to x. By (1) of Proposition 5.3, the set $\mathcal{P}$ is finite. Moreover since $p_u^*(x) \to y$ as $|u| \to +\infty$ (since u is a prefix of $p_u^*(x)$), it follows that $y \in \mathcal{P}$ and hence y is proximal to x. Whence by (2) of Proposition 5.3 we deduce that x and y are strongly coincident. □

Combining Theorems 5.1 and 3.13 we obtain

Corollary 5.4 ([7, Cor. 4.10]). *Let τ be an irreducible primitive substitution of Pisot type. Then a pair of fixed points x and y of τ are strongly coincident if and only if for each prefix u of y, the set $x\big|_u$ is an IP-set.*

6. Dumont–Thomas numeration systems

In the previous section we obtained a reformulation of the strong coincidence condition for irreducible primitive substitutions of Pisot type in terms of IP-sets and hence in terms of additive combinatorics (see Corollary 5.4). It is therefore natural to consider certain numeration systems defined by substitutions first introduced by Dumont and Thomas in [20, 21]. Since in the irreducible Pisot case, condition (4) in Theorem 5.1 implies the strong coincidence condition, these numeration systems may provide a new insight to the strong coincidence conjecture.

Let τ denote a substitution on a finite alphabet $\mathbb{A}$. For simplicity we assume that τ has at least one fixed point $x = x_0x_1x_2\cdots$ beginning in some letter $a \in \mathbb{A}$. The idea behind the numeration system is quite natural: every coordinate x_n of the fixed point x is in the image of τ of some coordinate x_m with $m \leq n$. More precisely, consider the least positive integer m such that $x_0x_1\cdots x_n$ is a prefix of $\tau(x_0x_1\cdots x_m)$. In this case we can write $x_0x_1\cdots x_n = \tau(x_0x_1\cdots x_{m-1})u_nx_n$ where u_nx_n is a prefix of $\tau(x_m)$. We now imagine a directed arc from x_m to x_n labeled u_n. In this way every coordinate x_n is the target of exactly one arc, and the source of $|\tau(x_n)|$-many arcs. It follows that for each n there is a unique path s from x_0 to x_n. Thus every natural number n may be represented by a finite sequence of labels u_i obtained by reading the labels along the path s in the direction from x_0 to x_n.

More formally, associated to τ is a directed graph $\mathcal{G}(\tau)$ defined as follows: the vertex set of $\mathcal{G}(\tau)$ is the set $\mathbb{A}$. Given any pair of vertices a, b we draw a directed edge from a to b labeled $u \in \mathbb{A}^*$ if ub is a prefix of $\tau(a)$. In other words, for every occurrence of b in $\tau(a)$ there is a directed edge from a to b labeled by the prefix (possibly empty) of $\tau(a)$ preceding the given occurrence of b. Figure 4 depicts the graph $\mathcal{G}(\tau)$ for the Fibonacci substitution $a \mapsto ab$, $b \mapsto a$.

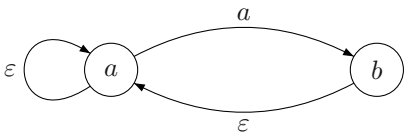

FIGURE 4. The Fibonacci automaton.

For simplicity, in case some letter b occurs multiple times in $\tau(a)$, we draw just one directed edge from a to b having multiple labels as described above. This is shown in Figure 5 in the case of the substitution $a \mapsto aba$, $b \mapsto baabb$.

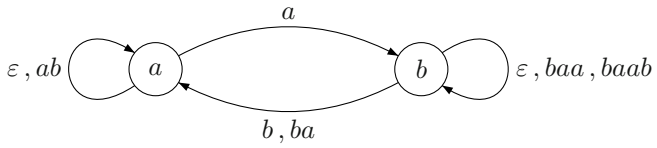

FIGURE 5. The automaton of $a \mapsto aba$, $b \mapsto baabb$.

Let $x = x_0x_1x_2\cdots$ denote the fixed point of τ beginning in a. Then the graph $\mathcal{G}(\tau)$ has a singleton loop based at a labeled with the empty word ε. We consider this to be the empty or 0th path at a. More generally by a path at $a \in \mathbb{A}$ we mean a finite sequence of edge labels $u_0u_1u_2\cdots u_n$ corresponding to a path in $\mathcal{G}(\tau)$ originating at vertex a with the condition that $u_0 \neq \varepsilon$ whenever the length of the path $n > 0$. For example in the case of the Fibonacci substitution, except for the path $s = \varepsilon$, each path is given by a word in $\{a, \varepsilon\}$ beginning in a and not containing the factor aa. For each path $s = u_0u_1u_2\cdots u_n$ set

$$\rho(s) = \tau^n(u_0)\tau^{n-1}(u_1)\tau^{n-2}(u_2)\cdots\tau(u_{n-1})u_n$$

and $\lambda(s) = |\rho(s)|$. In [20, 21] it is shown that for each path s at a, the word $\rho(s)$ is a prefix of the fixed point x at a and conversely for each prefix u of x there is a unique path s at a with $\rho(s) = u$. This correspondence defines a numeration system in which every natural number l is represented by the path $s = u_0u_1u_2\cdots u_n$ in $\mathcal{G}(\tau)$ from vertex a to vertex x_l corresponding to the prefix of length l of x, so that

$$(*) \quad l = \lambda(s) = |\tau^n(u_0)| + |\tau^{n-1}(u_1)| + |\tau^{n-2}(u_2)| + \cdots + |\tau(u_{n-1})| + |u_n|.$$

Generally by the numeration system one means the quantities $|\tau^n(u)|$ for all $n \geq 0$ and all proper prefixes u of the images under τ of the letters of $\mathbb{A}$. Then a *proper* representation of l in this numeration is an expression of the form (*) corresponding to a path $s = u_0u_1u_2\cdots u_n$ in $\mathcal{G}(\tau)$.

In the case of a uniform substitution of length k this corresponds to the usual base k-expansion of l. In the case of the Fibonacci substitution, each $u_n \in \{\varepsilon, a\}$ and $u_iu_{i+1} \neq aa$ for each $0 \leq i \leq n-1$. Thus this representation of l is the so-called Zeckendorff representation of l in which l is expressed as a sum of distinct Fibonacci numbers via the greedy algorithm (see [40]).

In general, this numeration system not only depends on the substitution τ but also on the choice of fixed point. For example for the substitution in Figure 2 the number 5 is represented by the path a, ba from vertex a or by the path b, ε from vertex b. In fact, $\tau(a)ba = ababa$ is the prefix of length 5 of $\tau^\infty(a)$ while $\tau(b)\varepsilon = baabb$ is the prefix of length 5 of $\tau^\infty(b)$.

An alternative reformulation is as follows: Given two distinct paths $s = u_0u_1u_2\cdots u_n$ and $t = v_0v_1v_2\cdots v_m$ both starting from the same vertex a, we write $s < t$ if either $n < m$ or if $n = m$ there exists $i \in \{0, 1, \ldots, n\}$ such that $u_j = v_j$ for $j < i$, and $|u_i| < |v_i|$. This defines a total order on the set of all paths starting from vertex a. In the case of the Fibonacci substitution, we list the paths

at a in increasing order

$$\varepsilon, a, a\varepsilon, a\varepsilon\varepsilon, a\varepsilon a, a\varepsilon\varepsilon\varepsilon, a\varepsilon\varepsilon a, a\varepsilon a\varepsilon, a\varepsilon\varepsilon\varepsilon\varepsilon, \ldots$$

Thus there is an order preserving correspondence between $0, 1, 2, 3, \ldots$ and the set of all paths at a ordered in increasing order.

While these numeration systems are very natural and simple to define, they are typically extremely difficult to work with in terms of addition and multiplication.

Let a and b be distinct vertices in $\mathcal{G}(\tau)$. We say a path s originating at a is *synchronizing* relative to b if there exists a path s' originating at b having the same terminal vertex as s and with $\lambda(s) = \lambda(s')$. From this point of view the strong coincidence conjecture implies that

$$\{\lambda(s) \,|\, s = \text{a synchronizing path relative to } b\}$$

is a thick set.

Let τ be a primitive substitution satisfying the strong coincidence condition. Suppose x and y are fixed points of τ beginning in a and b respectively. We will show that $x\big|_u$ is a central set for every prefix u of y. Since x and y are strongly coincident, we can write $x = scx'$, and $y = tcy'$ for some $s, t \in \mathbb{A}^+$, $c \in \mathbb{A}$, and $x', y' \in \mathbb{A}^\infty$ with $s \sim_{\mathrm{ab}} t$. By replacing τ by a sufficiently large power of τ, we can assume that

- sc is a prefix of $\tau(a)$,
- tc is a prefix of $\tau(b)$,
- b occurs in $\tau(c)$.

Thus in $\mathcal{G}(\tau)$ there is a directed edge from a to c labeled s, a directed edge from b to c labeled t, and a directed edge from c to b labeled r for some prefix r of $\tau(c)$. See Figure 6.

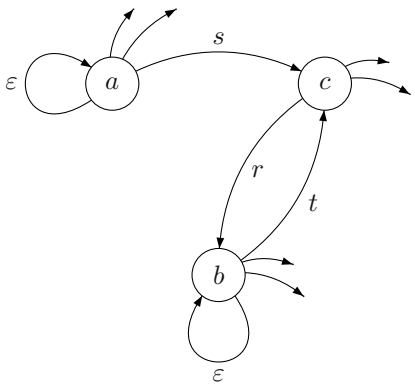

FIGURE 6. Vertices a, b, c of $\mathcal{G}(\tau)$.

We now define a sequence of paths $(p_i)_{i\geq 0}$ from a to b by

$$p_i = s, r, \underbrace{\varepsilon, \varepsilon, \ldots, \varepsilon}_{2i}.$$

Put $n_i = \lambda(p_i)$. Then clearly $\{n_i \,|\, i \geq 0\} \subseteq x\big|_b$. We now show that any finite sum of distinct elements from the set $\{n_i \,|\, i \geq 0\}$ is contained in $x\big|_b$. Set

$$q_i = t, r, \underbrace{\varepsilon, \varepsilon, \ldots, \varepsilon}_{2i}.$$

Then each q_i is a path from b to b and since s and t are Abelian equivalent it follows that $\lambda(p_i) = \lambda(q_i)$. Fix $k \geq 1$ and choose $i_1 < i_2 < \cdots < i_k$. Then

$$\begin{aligned}
\sum_{j=1}^{k} \lambda(p_{i_j}) &= \lambda(p_{i_k}) + \sum_{j=1}^{k-1} \lambda(p_{i_j}) \\
&= \lambda(p_{i_k}) + \sum_{j=1}^{k-1} \lambda(q_{i_j}) \\
&= |\tau^{2i_k+1}(s)| + |\tau^{2i_k}(r)| + \sum_{j=1}^{k-1} (|\tau^{2i_j+1}(t)| + |\tau^{2i_j}(r)|) \\
&= |\tau^{2i_k+1}(s)\tau^{2i_k}(r)\tau^{2i_{k-1}+1}(t)\tau^{2i_{k-1}}(r)\tau^{2i_{k-2}+1}(t)\tau^{2i_{k-2}}(r)\cdots \\
&\qquad\qquad \cdots \tau^{2i_1+1}(t)\tau^{2i_1}(r)|
\end{aligned}$$

which is represented by a path in $\mathcal{G}(\tau)$ from a to b and hence corresponds to an occurrence of b in x. This shows that $x\big|_b$ is an IP-set. It now follows from Theorem 3.13 that $x\big|_b$ is a central set. A similar argument applies for any prefix u of y by defining the paths p_i by

$$p_i = s, r, \underbrace{\varepsilon, \varepsilon, \ldots, \varepsilon}_{N_i}$$

with N_i sufficiently large.

7. Open problems

In this section we give a few open problems dealing with connections between word combinatorics and Ramsey theory. However, we first would like to draw attention to a list of nice open problems suggested by V. Bergelson and B. Rothschild in [11]. We begin with one of the problems presented there.

Although as we saw in Section 2.2 neither Σ_k nor Σ_k^∞ is partition regular, for each fixed k we could consider the set

$$\mathcal{R}^\infty(k) = \Big\{ A \subseteq \mathbb{N} \,|\, \text{whenever } r \in \mathbb{N} \text{ and } A = \bigcup_{i=0}^{r} A_i, \\ \exists\, 0 \leq i \leq r \text{ such that } A_i \text{ is } k^\infty\text{-summable} \Big\}$$

Then each $\mathcal{R}^\infty(k)$ is non-empty. In fact every IP-set belongs to $\mathcal{R}^\infty(k)$. The following is a difficult open question of Imre Leader:

Question 7.1 ([11, Question 8.1]). *Does there exist a member of $\mathcal{R}^\infty(2)$ which is not an IP-set?*

In general, the question of determining whether a given subset $A \subseteq \mathbb{N}$ is in $\mathcal{R}^\infty(k)$ or is an IP-set is typically quite difficult, even if for every A, either A or its complement belongs to $\mathcal{R}^\infty(k)$ or is an IP-set.

Now we need a few definitions:

Definition 7.2. Let $x \in \mathbb{A}^{\mathbb{N}}$, C a finite non-empty set and $\varphi : \mathcal{F}_x \to C$ a finite coloring of the factors of x. A factorization $x = V_0V_1V_2\cdots$ with $V_i \in \mathbb{A}^+$ is called

- *φ-monochromatic* if $\exists c \in C$ such that $\varphi(V_i) = c$ for all $i, j \geq 0$.
- *φ-super monochromatic* if $\exists c \in C$ such that $\varphi(V_iV_{i+1}\cdots V_{i+j}) = c$ for all $i, j \geq 0$.
- *φ-ultra monochromatic* if $\exists c \in C$ such that $\forall k \geq 1$, for all $n_1 < n_2 < \cdots < n_k$ and for all permutations σ of $\{1, 2, \ldots, k\}$ we have either $V_{n_{\sigma(1)}}V_{n_{\sigma(2)}}\cdots \cdots V_{n_{\sigma(k)}} \notin \mathcal{F}_x$ or $\varphi(V_{n_{\sigma(1)}}V_{n_{\sigma(2)}}\cdots V_{n_{\sigma(k)}}) = c$.

For instance, if $x = u^\omega$ is periodic with $u \in \mathbb{A}^+$, then the factorization $x = u \cdot u \cdot u \cdots$ is φ-monochromatic for any finite coloring $\varphi : \mathcal{F}_x \to C$. In general this factorization need not be φ-super monochromatic.

In [19] we conjectured that if x is not periodic then there exists a finite coloring $\varphi : \mathcal{F}_x \to C$ such that no factorization of x is φ-monochromatic. The conjecture is verified for all non-uniformly recurrent words and various classes of uniformly recurrent words including Sturmian words (see [19]). What is immediate to see is that if x is not periodic, then there exists a finite coloring $\varphi : \mathcal{F}_x \to C$ such that no factorization of x is φ-super monochromatic. In fact, it suffices to take $\varphi : \mathcal{F}_x \to \{0, 1\}$ defined by $\varphi(u) = 0$ if u is a prefix of x and $\varphi(u) = 1$ otherwise. Suppose to the contrary that x admits a φ-super monochromatic factorization $x = V_0V_1V_2\cdots$. Then since V_0 is a prefix of x, it follows that $V_1V_2\cdots V_k$ is a prefix of x for each $k \geq 1$. Thus $x = V_0x$, whence x is periodic ($x = V_0V_0V_0\cdots$), a contradiction. However, it could be shown that for any word x and any finite coloring $\varphi : \mathrm{Fact}(x) \to C$, some suffix of x admits a φ-super monochromatic factorization.

Question 7.3. *Prove or Disprove: Let $x \in \mathbb{A}^{\mathbb{N}}$ be non-periodic. Then there exists a finite coloring $\varphi : \mathcal{F}_x \to C$ such that x does not admit a φ-monochromatic factorization.*

Question 7.4. *Prove or Disprove: Let $x \in \mathbb{A}^{\mathbb{N}}$ and $\varphi : \mathcal{F}_x \to C$ a finite coloring. Then there exists a suffix x' of x which admits a φ-ultra monochromatic factorization.*

It can be shown that as a consequence of Hindman's theorem, Question 7.4 has an affirmative answer in case x is ultimately periodic. It can also be shown

that as a consequence of Ramsey's infinite theorem, for any word $x \in \mathbb{A}^{\mathbb{N}}$ and any finite coloring $\varphi : \mathcal{F}_x \to C$, there exists a suffix x' of x which admits a φ-super monochromatic factorization. Moreover, an affirmative answer to Question 7.4 implies both Hindman's theorem and Ramsey's theorem.

References

[1] J.-P. Allouche and J. Shallit, *The ubiquitous Prouhet–Thue–Morse sequence, in: Sequences and their applications*, Proceedings of SETA'98, C. Ding, T. Helleseth and H. Niederreiter (eds.) (1999), Springer Verlag, 1–16.

[2] J.-P. Allouche and J. Shallit, *Automatic sequences: Theory, applications, generalizations*, Cambridge University Press, 2003.

[3] M. Barge and B. Diamond, *Coincidence for substitutions of Pisot type*, Bull. Soc. Math. France **130** (2002), 619–626.

[4] M. Barge and B. Diamond, *Proximality in Pisot tiling spaces*, Fund. Math. **194** (2007), 191–238.

[5] M. Barge and J. Kwapisz, *Geometric theory of unimodular Pisot substitutions*, Amer. J. Math. **128** (2006), 1219–1282.

[6] M. Barge and J. Kwapisz, *Elements of the theory of unimodular Pisot substitutions with an application to β-shifts*, Algebraic and Topological Dynamics, 89–99, Contemp. Math. **385**, Amer. Math. Soc., Providence, RI, 2005.

[7] M. Barge and L.Q. Zamboni, *Central sets and substitutive dynamical systems*, Adv. Math. **248** (2013), pp. 308–323.

[8] V. Bergelson, *Minimal idempotents and ergodic Ramsey theory*, Topics in dynamics and ergodic theory, 8–39, London Math. Soc. Lecture Note Ser., **310**, Cambridge Univ. Press, Cambridge, 2003.

[9] V. Bergelson and N. Hindman, *Nonmetrizable topological dynamics and Ramsey Theory* Trans. Amer. Math. Soc. **320** (1990), pp. 293–320.

[10] V. Bergelson, N. Hindman and D. Strauss, *Strongly central sets and sets of polynomial returns mod* 1, Proc. Amer. Math. Soc., **140** (2012), 2671–2686. to appear.

[11] V. Bergelson and B. Rothschild, *A selection of open problems*, Topology Appl. **156** (2009), 2674–2681.

[12] V. Bergelson and Y. Son, Personal communication.

[13] A. Blass, *Ultrafilters: where topological dynamics = algebra = combinatorics*, Topology Proc. **18** (1993), pp. 33–56.

[14] M. Bucci, N. Hindman, S. Puzynina and L.Q. Zamboni, *On additive properties of sets defined by the Thue–Morse word*, J. Comb. Theory, Series A **120** (2013), 1235–1245.

[15] M. Bucci, S. Puzynina and L.Q. Zamboni, *Central sets generated by uniformly recurrent words*, Ergodic Theory & Dynam. Sys., (2013), in press.

[16] D. De, N. Hindman and D. Strauss, *A New and Stronger Central Sets Theorem*, Fundamenta Mathematicae **199** (2008), p. 155–175.

[17] A. de Luca, *Sturmian words: structure, combinatorics, and their arithmetics*. Theoret. Comput. Sci. **183** (1997), 45–82.

[18] F.M. Dekking, *The spectrum of dynamical systems arising from substitutions of constant length*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, **41** (1977/1978), 221–239.

[19] A. de Luca, E. Pribavkina, L.Q. Zamboni, *A coloring problem for infinite words*, J. Combinatorial Theory (Series A) **125** (2014), 306–332.

[20] J.-M. Dumont and A. Thomas, *Systèmes de numération et fonctions fractales relatifs aux substitutions*, Theoret. Comput. Sci., **65** (2) (1989), 153–169.

[21] J.-M. Dumont and A. Thomas, *Digital sum moments and substitutions*, Acta Arith., **64** (1993), 205–225.

[22] S. Dworkin, *Spectral theory and X-ray diffraction*, J. Math. Phys. **34** (1993), 2965–2967.

[23] R. Ellis, *Distal transformation groups*, Pac. J. Math. **8** (1958), p. 401–405.

[24] H. Furstenberg, *Recurrence in ergodic theory and combinatorial number theory*, Princeton University Press, Princeton, 1981.

[25] N. Hindman, *Finite sums of sequences within cells of a partition of* $\mathbb{N}$, J. Combinatorial Theory (Series A) **17** (1974), p. 1–11.

[26] N. Hindman, I. Leader and D. Strauss, *Infinite partition regular matrices: solutions in central sets*, Trans. Amer. Math. Soc. **355** (2003), 1213–1235.

[27] N. Hindman and D. Strauss, *Algebra in the Stone–Čech compactification: theory and applications*, 2nd edition, Walter de Gruyter & Co., Berlin, 2012.

[28] N. Hindman and D. Strauss, *A simple characterization of sets satisfying the Central Sets Theorem*, New York J. Math **15** (2009), p. 405–413.

[29] J.-Y. Lee, R. Moody and B. Solomyak, *Consequences of Pure Point Diffraction Spectra for Multiset Substitution Systems*, Discrete Comp. Geom. **29** (2003), 525–560.

[30] M. Morse, *Recurrent geodesics on a surface of negative curvature*, Trans. Amer. Math. Soc. **22** (1921), p. 84–100.

[31] M. Morse and G. Hedlund, *Symbolic dynamics*, Amer. J. Math. **60** (1938), p. 815–866.

[32] M. Morse and G. Hedlund, *Symbolic dynamics II: Sturmian sequences*, Amer. J. Math. **62** (1940), p. 1–42.

[33] R. Rado, *Studien zur Kombinatorik*, Math. Zeit. **36** (1933), 242–280.

[34] F. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. **30** (1930), 264–286.

[35] G. Rauzy, *Nombres algébriques et substitutions*, Bull. Soc. Math. France **110** (1982), 147–178.

[36] G. Rauzy, *Rotations sur les groupes, nombres algébriques, et substitutions*, in Séminaire de Théorie des Nombres (Talence, 1987–1988), Université de Bordeaux 1, exposé n. 21 (1988).

[37] I. Schur, *Über die Kongruenz* $x^m + y^m = z^m (\mod p)$, Jahresbericht der Deutschen Math. Verein. **25** (1916), 114–117.

[38] A. Thue, *Über unendliche Zeichenreihen*, Norske Vid. Selsk. Skr. I Math-Nat. Kl. **7** (1906), 1–22.

[39] B. van der Waerden, *Beweis einer Baudetschen Vermutung*, Nieuw Arch. Wiskunde **19** (1927), 212–216.

[40] E. Zeckendorff, *Représentation des nombres naturels par une somme de nombres de Fibonacci ou de nombres de Lucas*, Bull. Soc. Royale Sci. Liège, **42** (1972), pp. 179–182.

Svetlana Puzynina
Sobolev Institute of Mathematics
4 Acad. Koptyug avenue
630090 Novosibirsk, Russia

and

University of Turku
Department of Mathematics
20014 Turun yliopisto, Finland
e-mail: svepuz@utu.fi

Luca Q. Zamboni
Université Claude Bernard Lyon 1
43 boulevard du 11 novembre 1918
F-69622 Villeurbanne cedex, France

and

University of Turku
Department of Mathematics
20014 Turun yliopisto, Finland
e-mail: lupastis@gmail.com

Progress in Mathematics, Vol. 309, 405–428
© 2015 Springer Basel

Delone Sets and Material Science: a Program

Jean V. Bellissard

Abstract. These notes are proposing a program liable to provide physicists working in material science, especially metallic liquids and glasses, the mathematical tools they need to build an atomic scale theory of Continuous Mechanics including plasticity, fluidity and, hopefully, fractures. Using the long list of datas and numerical simulations accumulated during the last forty years, physicists have identified a new class of degrees of freedom, besides the elastic ones, which will be called *anankeons* here [7]. They are dominant in the liquid phase and they explain the properties related to plastic deformations of the solid phase. It is advocated that Delone sets provide a natural frame within which such a theory can be expressed. The use of Voronoi tiling and its dual construction, called the Delaunay triangulation, gives a discretization of the data. The concept of Pachner move or Delaunay flips permits to describe very precisely what the anankeons are. A partition of the configuration space into contiguity domains leads to a graph on which a Markov process can be built to describe the anakeon dynamics. At last, a speculative Section is giving an attempt to describe the Continuous Mechanics of a condensed material in terms of a Noncommutative Geometry of the configuration space.

Mathematics Subject Classification (2010). 74A02, 70F02, 82-02.

Keywords. Bulk metallic glasses, Delone sets, Pachner move, STZ, anankeons.

1. Introduction

This article is an attempt to explain how the mathematical tools used in the study of tiling spaces can be used to describe very concrete materials like liquids and glasses. The text will then have a part describing the situation viewed from the eye of the physicists or the material scientist, and another part consisting in describing the mathematical material. As a consequence, the reader who is not an expert in Physics or in Statistical Mechanics, may need some time to adjust

Work supported in part by NSF grant DMS 1160962.

and understand the content. The author apologizes for demanding so much from the reader. However, developing a mathematical tool liable to describe the real world is never an easy task because it always includes several important aspects requiring to use pieces of Mathematics that have usually little to do with each other.

1.1. A program

Describing a material in a condensed phase, solid or liquid, from the atomic scale on is still a challenge today. These notes provide a guideline liable to fill this gap in the future. The main new ingredient is the concept of *anankeon* [1, 7]. This is the dominant degree of freedom for liquids, while anankeons play a crucial role in solids to account for their plasticity. As will be seen, this concept of anankeon fits with a description of atomic configurations in terms of Delone sets, Voronoi tiling, Delaunay triangulation. Anankeon then appear to be any of the individual simplex of this triangulation. Changes in the local triangulations, known as *Pachner moves* or *Delaunay flips*, give a way to describe the anankeon dynamics. However the latter is essentially unpredictable and has to be given through a Markov process. Numerical simulations, based on molecular dynamics are suggesting various phenomenological forms for this Markov process.

In the last section of these notes a list of speculations will be presented in form of a program. The author believes, indeed, that the mechanical properties of a condensed material should have a geometrical interpretation. He will argue that the Geometry is not in the usual space but in the *Transversal* direction of the *Hull* of the set of atomic configurations. The argument will be given along the following lines. The generator of a Markov process is usually considered as a Dirichlet form acting on a Hilbert space [51, 52, 59]. In the present situation, however, it is expected to be rather the generator of a Markov semigroup on an abelian C^*-algebra $\mathcal{A}$ attached to *Transversal*. The invariance of $\mathcal{A}$ by the Markov semigroup is the analog of the *Feller property* in the theory of Markov Processes. It will be argued that, in such a case, it should be possible to define a metric on the tiling space in a way similar to the constructions of Kantorovich [61, 62] or Connes [55]. It is then expected to provide the tiling space with a Noncommutative analog of a Riemannian manifold. The main issue discussed at the end is the definition of the *curvature*. This concept is still ill defined in Noncommutative Geometry. However, the curvature is expected to be related to the mechanical properties of the material, through some kind of Einstein equation:

$$\textbf{Curvature} = \textbf{Stress}$$

1.2. Some physics background

The usual approach to solids and liquids is to distinguished between scales: length scales, time scale, energy scales. At macroscopic length scales either the Elasticity Theory for solids (see for instance [14, 8]), or the Navier–Stokes equation for fluids, give the standard description. At a lower scale, a kinetic theory approach, using the Bolztmann equation or the BBGKY hierarchy, is also considered (see for instance

[11]). While kinetic theories have been successful for gas, they are more difficult to manipulate for liquids. Since the late nineties two major theoretical contributions were made to improve the description of solids at large scale. The first is due to a group of physicists led by J. Langer, and is called *the STZ Theory*, where STZ stands for *shear-transformation zones* [9, 10, 15, 16, 17]. It proved to be very efficient in describing the plasticity region in the strain-stress curve, to describe also the evolution of fractures [20, 19] in glassy systems. The other one, more phenomenological, is called *Peridynamics Theory* and was initiated by Silling [21, 22, 23]. It became recently important in particular to describe solid under extreme stress, like explosives.

The discovery of a new class of materials, during the nineties, called *bulk metallic glasses* (BMG), like the vitalloys containing mostly copper and zirconium, has been an experimental and a technological breakthrough [12, 13]: these materials have exceptional mechanical properties, a very high limit of elasticity and an exceptional resistance to rupture. They are now used in various devices including sport material (i-Phone6, golf clubs, ...) and medical tools. The development of these materials has triggered more research, has attracted funding and therefore it provides a window of opportunity for theoreticians and experimentalists alike to test ideas and models liable to give access to a better understanding of both glass and liquid phases.

1.3. Lessons from available data

Several experts in the field of BMG have expressed the wish to develop a more universal modeling valid from the atomic scale on with the help of mathematicians. There are several layers of difficulties in this problem. In the fifties, Bernal used a Voronoi construction to describe the atomic configurations of such materials [3]. In the early eighties, the *local cluster theory* provided a more detailed description of the local environment of a typical atom. This was initiated by T. Egami *et al.* [6] and was completed by Miracle [18] during the last decade. The main experimental tool to analyze the results of this approach is the diffraction spectrum (X-ray, electrons or neutrons), the Fourier transform of which permits to compute the *pair distribution function* (PDF).

A *liquid phase approach* has also been advocated for a long time, in order to identify the most relevant degrees of freedom. The *heat capacity* is an observable easy to measure in experiments. It provides some informations about how the energy is stored within various degrees of freedom. The behavior of the heat capacity of a typical disordered material (see Figure 1) reveals a low temperature regime in the solid phase (glass), a saturation at high temperature (liquid phase) leading to a similarity with the law of Dulong–Petit and a transition region (liquid-glass transition) showing an increase of the heat capacity which looks like a phase transition without a singularity. The *Law of Dulong–Petit* predicts that the heat capacity per unit volume of a solid saturates at high temperature and is given by

$$C_V = Mk_B\,, \qquad M = \text{number of degrees of freedom}\,, \tag{1}$$

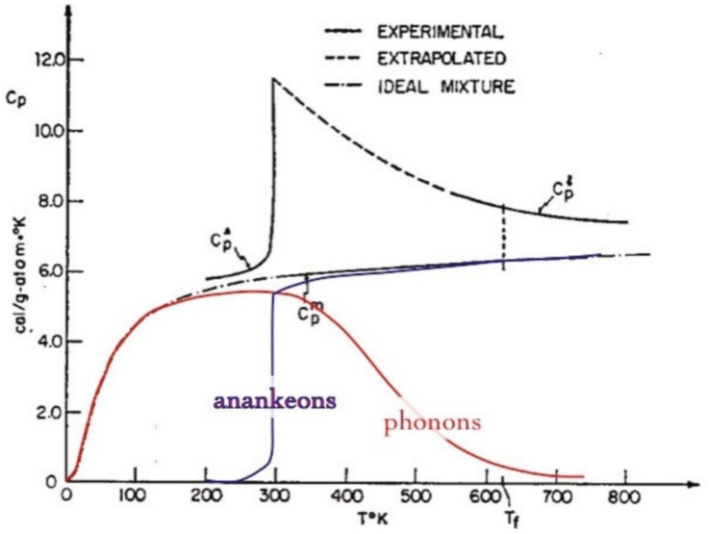

FIGURE 1. Heat capacity as a function of temperature for an alloy made of Gold, Germanium and Silicon signaling a glass-liquid transition [4]. In color a possible contribution of phonons and anankeons.

and k_B is the Boltzmann constant. Einstein assumed that the main degrees of freedom in a solid where provided by *phonons*, namely the quantized version of acoustic waves, leading to $M = 3$. However, it is known that phonons are mainly damped in liquids so that the experimental result showed in Figure 1 requires an explanation. The anakeon theory might give an interpretation of these observations (see Section 3.1 below).

The second layer of difficulty is the time dependence of the response to mechanical strain. For example diving into water by jumping from the edge of a swimming pool is possible because the water moves fast enough to permit the diver's body to enter smoothly. However, firing a bullet in water results in the bullet being completely smashed like if water was a piece of concrete. In the former case, the speed of the diver is low, permitting the water to react fast enough, while in the latter case, the speed of the bullet is so high that the water is too slow to adjust around the bullet and behave like a very hard solid. In more scientific terms, the response to shear involves a time scale that is usually expressed through the *viscosity* of the medium. In particular, the main difference between liquid and glass is not seen in the equilibrium parameters like density or heat capacity. It is mainly seen in a sudden increase of the viscosity. Namely the time scale characterizing the response to shear increases by several orders of magnitude within a very small interval of temperature at the transition. The challenge in this case is to provide a theory liable to describe the *dynamics* of the atomic configurations.

1.4. Degrees of freedom

Over the years, the work of Egami and his collaborators has permit to identify two relevant classes of degrees of freedom at the atomic scale.

1. the *vibration* degrees of freedom, the one at work in the theory of elasticity, at least for long wave length or the ones with short wave-lengths that can eventually be quantized and which has been called *phonon* by A. Einstein in 1908. In the glass phase, phonon seem to give the dominant contribution. However in the liquid phase the high frequency phonons are damped, with a very short lifetime.
2. the topological degrees of freedom, representing the unpredictable sudden change in the local configuration of atoms due to the local stress and the thermal motion. These degrees of freedom are newcomers in the theoretical landscape and will be called *anankeon* here [1]. They provide an explanation to the Dulong–Petit behavior of the heat capacity [5] (see Figure 1), namely the saturation at large temperature. They just freeze at the liquid-glass transition. They reappear only if an external shear is applied as the atomic scale version of the STZ proposed by Langer *et al.*.

While the dynamics of phonons is well understood, including nonlinearity if necessary, the dynamics of anankeon is not. A model for this dynamics, using a Markov process, will be proposed. The construction of such a process uses the Boltzmann-Gibbs factor for the equilibrium dynamics as well as the Arrhenius law for the evaluation of waiting times. Recent numerical simulations made by T. Egami and his group of collaborators, give a hint about the actual magnitude of the various parameters, energy, length and time scales involved in such a modeling.

2. Why Delone sets?

2.1. Interatomic potential

Even though atoms, especially in transition metals like copper or zirconium, have a complicated structure with many electrons around the nucleus, it is a convenient and usual approximation to consider them as individual particles interacting through a pair-potential. This potential is usually a function of the relative position of the atoms having the following properties: (i) it is repulsive at short distance and (ii) attractive at large distance. In addition, for most isotropic materials, it can be assumed that this potential is rotation invariant, meaning that it is a function of the interatomic distance r. Examples used in numerical simulations, using molecular dynamic, for instance, are the Lennard–Jones potential, valid mostly for systems interacting with the van der Waals force, such as rare gases, the Johnson one which is short range and exhibit oscillations or potential computed *ab initio* for metallic alloys (see Figure 2). Apparently the results of simulations are not strongly affected by the nature of the potential. The origin of such potential from the first principle will not be discussed here. However they provide some concrete length scales: (i) a short length scale σ such that for $r < \sigma$ then $V(r) > 0$, (ii) the

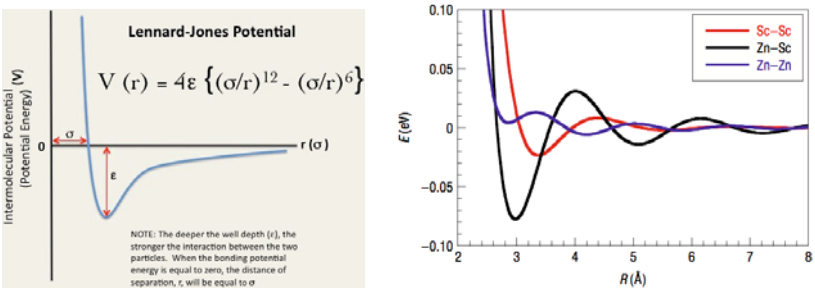

FIGURE 2. *Left*: a representation of the Lennard–Jones potential; *Right*: an effective interatomic potential computed numerically for an $Zn - Sc$ alloy.

minimum potential length a. It is expected that atomic configurations with two atoms closer than σ are rare. In addition, atomic configurations with a vacancy bigger than a are unstable *w.r.t.* shear (see Figure 4 and a discussion in [25]), giving a limit for the typical size of a vacancy.

2.2. Quasi-discrete measures

A configuration of atoms can be represented by the positions of its atomic nuclei. Namely, if $\mathcal{A}$ denotes the set of atomic species present in the material, the position of atoms of the species a make up a discrete sets of points $\mathcal{L}_a \subset \mathbb{R}^d$ where $a \in \mathcal{A}$. Such a set is *topologically closed*, meaning that it has no accumulation point. It is mathematically better represented by the corresponding sum of Dirac measures δ_x located at the position of the atomic nuclei, namely by $\nu_a = \sum_{x\in\mathcal{L}_a} \delta_x$. Such a measure are called *discrete* and can be characterized by the following properties [2]:

(i) it is a Radon measure on $\mathbb{R}^d$,
(ii) given any ball $B \subset \mathbb{R}^d$ the measure $\nu_a(B)$ of this ball is an integer for each $a \in \mathcal{A}$.
(iii) the measure of a point $\nu_a\{x\}$ is either 0 or 1.

What is a Radon measure? It is a way to evaluate weighted integral of functions. Which functions? Any complex-valued function f defined on the physical space $\mathbb{R}^d$ (i) which is continuous and (ii) which vanishes outside of some bounded ball (namely F has *compact support*). The space of such functions is denoted by $\mathcal{C}_c(\mathbb{R}^d)$. Then the integral

$$\mu(f) = \int_{\mathbb{R}^d} f(x)\mu(dx)$$

has the following properties: (i) $\mu(f)$ is linear in f, (ii) if $(f_n)_{n=1}^{\infty}$ is a sequence of function in $\mathcal{C}_c(\mathbb{R}^d)$ which are all supported a common ball B and converging uniformly on B to a function f (that is to say $\lim_{n\to\infty} \sup_{x\in B} |f(x) - f_n(x)| = 0$),

then $\mu(f) = \lim_{n\to\infty} \mu(f_n)$. In the previous case

$$\nu_a(f) = \int_{\mathbb{R}^d} f(x)\nu_a(dx) \stackrel{def}{=} \sum_{x\in\mathcal{L}_a} f(x)\,.$$

The set of Radon measures is equipped with its weak*-topology, meaning that a sequence $(\nu_n)_{n\in\mathbb{N}}$ of such measures converges to a measure ν if and only if $\nu_n(f) \to \nu(f)$ for any complex-valued continuous function f on $\mathbb{R}^d$ with compact support. As it turns out, the set of discrete measures described above is not weak*-closed. Its closure is made of measures of the form $\nu = \sum_{x\in\mathcal{L}} n_x\delta_x$, where $\mathcal{L} \subset \mathbb{R}^d$ is closed and discrete, while the n_x's are integers. Namely it describes configurations of atoms in which a finite number of atoms are allowed to occupy the same position. This weak*-closure is denoted by $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$, where QD stands for *quasi-discrete*. It is a *Polish space* [2]: namely, its topology can be described through a metric for which it is complete, while such a metric is neither unique nor canonical. Polish spaces allow for a good theory of probability [36]. In addition, in the present case, given a family of pair-potentials $(V_{ab})_{a,b\in\mathcal{A}}$ of the previous type and a bounded open set U, it becomes possible to compute the local potential energy of a *discrete measure* $\nu = (\nu_a)_{a\in\mathcal{A}}$ as follows

$$E_{U,r,R}(\nu) = \sum_{a,b\in\mathcal{A}} \sum_{r<|x-y|<R\,;\,x\in U} \nu_a\{x\}\ \nu_b\{y\}\ V_{ab}(|x-y|)\,,$$

$$E_U(\nu) = \sup_{r>0,R>0} E_{U,r}(\nu)\,.$$

It can be shown that $E_{U,r,R}$ is a weak*-continuous function of ν as long as $r > 0$ and $R < \infty$. Consequently E_U is lower semicontinuous (in particular configurations with more than one atom at the same point inside U have infinite local energy). The previous energy can be decomposed into the energy internal to U, corresponding to the pairs x, y in U and the interaction between U and its complement $U^c = \mathbb{R}^d \setminus U$. Hence

$$E_U(\nu) = E_U^{\mathrm{int}}(\nu) + E_{U;U^c}(\nu)\,.$$

A real- or complex-valued function over $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ will be called *cylindrical* if it is continuous and if there is a bounded open set U such that F depends only upon the restriction of ν to U. Then U will be set to *support* F.

Definition 1. A configuration of atoms, namely a point in QD, is Delone if there are $0 < r < R < \infty$ such that in any ball of radius r there is at most one atom, while in any ball of radius R there is at least one atom. The family of Delone sets corresponding to the parameters r, R is denoted by $\mathrm{Del}_{r,R}(\mathbb{R}^d)$.

Delone Set Hypothesis: *The only physically relevant configuration are Delone sets.*

This hypothesis is justified by two observations. First two atoms are unlikely to get too close to each other, due to the strong repulsion at short distance. Second, in a condensed phase where the atomic density is high, the attracting part of the potential makes unlikely the occurrence of some large region void of atoms. That

such an hypothesis is reasonable will be discussed in the forthcoming sections. The following result is fundamental.

Proposition 1 (see [2]). *The space* $\mathrm{Del}_{r,R}(\mathbb{R}^d)$ *is a weak*-compact and translation invariant subspace of* $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$. *Any translation acts as an homeomorphism. In addition, given any Delone measure* ν *and any open ball* B *with diameter larger than* $2R$, *the number of atoms in* B *is bounded by*

$$\left(\frac{\mathrm{diam}(B)}{2R}-1\right)^d \leq \nu(B) \leq \left(1+\frac{\mathrm{diam}(B)}{2r}\right)^d .$$

Definition 2. Given $\nu \in \mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ its Hull is the weak*-closure of its orbit under the translation group. In particular, if ν is (r,R)-Delone Hull$(\nu) \subset \mathrm{Del}_{r,R}(\mathbb{R}^d)$ is compact and the translation group acts by homeomorphism.

2.3. Thermal equilibrium

Using a pair potential of the previous type, it becomes possible to describe the thermal equilibrium in a rigorous way using a local Gibbs distribution. Let F be a cylindrical function supported by U. Then the *thermal average* of F in U is usually defined by [27, 28]

$$\mathbb{E}_U(F)(\nu) = \frac{1}{Z_U(\beta,\mu;\nu)} \sum_{N_a=0}^{\infty} \prod_{a\in\mathcal{A}} \frac{e^{\beta\mu_a N_a}}{N_a!} \times \int_{x_{k_a}\in U} \prod_{a\in\mathcal{A}} \prod_{k_a=1}^{N_a} dx_{k_a}\; e^{-\beta(E_U^{\mathrm{int}}(x)+E_{U;U^c}(x,\nu))}\; F(x)$$

where $\beta = 1/k_B T$ denotes the inverse temperature measure in units of energy, $\mu = (\mu_a)_{a\in\mathcal{A}}$ denotes the family of *chemical potentials* associated with each atomic species, $x = ((x_{k_a})_{k_a=0}^{N_a})_{a\in\mathcal{A}}$ denotes the (random) positions of the atoms of each species inside U, N_a denotes the (random) number of atoms of species a inside U, while $\nu \in \mathrm{QD}_{\mathcal{A}}$ fixes the position of the atoms outside of U. As it is usual, the chemical potential is fixed by the atomic densities of each atomic species at a given temperature. Hence $\mathbb{E}_U$ defines a probability measure on $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ called (local or U-) *Gibbs state*. It is worth noticing that the kinetic contribution to the Gibbs state factorizes out and can be ignored here. It is also important to realize that since the configurations having two atoms or more on top of each other inside U have an infinite energy, they have zero probability *w.r.t.* the local Gibbs state.

In principle, the infinite volume limit is rigorously controlled through the so-called *Dobrushin–Lanford–Ruelle equations* (DLR) (see [26, 30, 27] for the DLR equations and [31, 27, 28] for the formalism with classical particle systems). Namely a *Gibbs state is a probability measure* $\mathbb{P}$ *over* $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ *admitting* E_U *as conditional probability when conditioned by the configurations outside* U. In particular such a probability is *locally absolutely continuous*. It is known that, under certain conditions on the pair-potentials, the set $\mathfrak{G}$ of Gibbs states is a *Choquet simplex*,

namely it is *compact* and convex. In particular $\mathfrak{G}$ admits extremal points and any Gibbs state is a *unique* convex combination of extremal points. The same property occurs for the subset $\mathfrak{G}_t$ of *translation invariant* Gibbs states. As it turns out, the extremal states in $\mathfrak{G}_t$ have two important properties:

(i) they are ergodic under the translation group,
(ii) they correspond to the pure homogeneous phases.

2.4. The ergodicity paradox

The property of extreme translation invariant Gibbs states of being ergodic under the translation group has several important consequences summarized as follows.

Proposition 2 (see [2]). *Let $\mathbb{P}$ be a translation invariant ergodic measure on $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$. Then*

(i) *there is a weakly closed subset $\Omega \subset \mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ such that for $\mathbb{P}$-almost every ν the Hull of ν coincides with Ω;*
(ii) *If $\mathbb{P}$ is supported by the space of Delone measures, then there are $0 < r \leq R < \infty$ such that $\mathbb{P}\{\mathrm{Del}_{r,R}(\mathbb{R}^d)\} = 1$ and $\mathbb{P}\{\mathrm{Del}_{r',R'}(\mathbb{R}^d)\} = 0$ if $r' \leq r, R' \geq R$ and $(r', R') \neq (r, R)$.*

The main problem with this approach comes from the following remark: (i) the set $\mathrm{Del}(\mathbb{R}^d)$ of all Delone sets is translation invariant, (ii) it is a Borel subset of $\mathrm{QD}_{\mathcal{A}}(\mathbb{R}^d)$ as a countable union of compact sets. Hence the ergodicity of $\mathbb{P}$ implies that the set of Delone sets has either probability zero or probability one. However it is quite standard to prove that at *positive temperature*, the set of configurations having either an arbitrary big hole somewhere or an arbitrary number of particles in a small ball somewhere else, has positive density. This can be done by evaluating the probability of a big hole in a finite volume. This density is very small, but it is not zero. Hence at positive temperature $\mathbb{P}(\mathrm{Del}(\mathbb{R}^d)) = 0$ whenever $\mathbb{P}$ is a translation invariant pure phase. Consequently the *Delone set Hypothesis is mathematically incorrect*.

There are two ways to reconcile Ergodicity with the Delone Hypothesis. The first one is to remark that in the limit of zero temperature only the minimal energy configuration (groundstates) matter. If a big hole occurs, it is energetically more favorable to fill it. Similarly an accumulation of particles in a small ball is energetically unfavorable. This is expressed by the following result [24].

Theorem 1 (see [24]). *Let the two-body potentials V_{ab} be hard-core, namely there is $r > 0$ such that $V_{ab}(x) = +\infty$ if $|x| \leq r$. Then, whenever the chemical potentials belong to a range allowed by the densities of each species, there is $R > r$ such that each ground-state configuration belongs to $\mathrm{Del}_{r,R}(\mathbb{R}^d)$.*

In the previous statement, the hard-core hypothesis is artificial because it puts by hand a property that should be spontaneously satisfied by the system. So the following problem is expected to have a positive answer.

Problem 1. Prove or disprove that there is a physically relevant class of rotation invariant pair potentials such that, for a high enough density of atoms, given the concentration of each atomic species, any zero temperature limit point of translation invariant Gibbs states gives probability one to $\mathrm{Del}(\mathbb{R}^d)$. □

The second way to reconcile Ergodicity with the Delone Hypothesis is to appeal to the *persistence theory* [29]. Namely rare events have a very short lifetime. This is why large vacancies are not observed in liquids and solids [25]. Similarly concentration of many atoms in a small ball is essentially impossible to observe. If, however, such extreme configurations were easy to observe, then it could be possible, for instance, to force the fusion of nuclei, namely to allow for *cold fusion* experiments. Hence there is a problem with the Thermal Equilibrium approach: no time scale is present. Introducing the time evolution in the problem is not easy, but there should be a way to do so in order to reconcile the observations with the mathematical approach.

Problem 2. Is it possible to develop a mathematically rigorous time dependent approach in order to reconcile the Delone Hypothesis with the translation Ergodicity of Gibbs states? More generally, is it possible to formulate and to develop a finite time approximation of an ergodic dynamical system permitting to neglect rare event with too short lifetime? □

3. The anankeon theory

This section provides few arguments, found by physicists, leading to the concept of *anankeon*.

3.1. Glasses, frustration, local stress

It is a known fact of practice that most materials, in particular metals, have a tendency to crystallize. Hence producing glasses with metals has been a real challenge (see [37] for a review). The earliest technique was melt quenching, namely an ultrafast cooling with cooling rate of $10^5 - 10^6\ K/s$. This technique limits sample to thin ribbons. Other techniques were developed later. Eventually a compound with low cooling rate was produced by Inoue *et al.* (see [12]). Later the class of *vitalloy* was produced by Johnson *et al.* (see [13]) and is nowadays the most studied class of bulk metallic glasses. Vitalloy can be casted in large samples of several centimeters. It was quickly realized that these new materials had exceptional mechanical properties, boosting the subject in the front stage of material science.

Thanks to the theoretical modeling of *spin-glass* the concept of *frustration* emerged as a fundamental aspect of glassy states [35]. It was slowly understood that such a frustration occurs in metallic glasses whenever several atomic species with *different atomic radii* where mixed together in such proportions as to force atoms of different size to be close to each other. In this way the number of small clusters that can be formed in the melt is getting larger [6, 18] and crystallization becomes unlikely.

Another ingredient that should be included, advocated for a long time by Egami *et al.*, is the local stress exercised by neighboring atoms on a single one. For indeed, as can be shown by a Thomas–Fermi approximation [32, 34], atoms are not exactly hard ball. They can be squeezed a little, however big is their resistance to pressure. This leads to associating with each atom a *stress tensor* σ_i, by adding the contributions of forces acted upon by the neighboring atoms. An important observation, made through numerical simulations using molecular dynamics, is the following:

Egami-Srolovitz Principle: *In the liquid phase, the atomic stress tensors behave like independent identically distributed random Gaussian variables* [33]

More precisely, a stress tensor σ is a 3×3 real symmetric matrix. It is usually decomposed into its trace part $p = (1/3)\mathrm{Tr}\,(\sigma)$, called the *pressure*, and its traceless part $\widehat{\tau} = \sigma - p$ called the *deviatoric stress*. Then the parameter τ, defined by $\tau^2 = (3/2)\mathrm{Tr}\,(\widehat{\sigma}^2)$, is called the *von Mises stress*. The atomic stress distribution is given by [33] the six-dimensional integral

$$\mathrm{Prob}\{\sigma_i \in A\} = \frac{1}{Z}\int_A e^{-\beta(p^2/2B+\tau^2/2G)}\, d\sigma\,.$$

where $A \subset M_3^s(\mathbb{R})$ is any Borel set, $d\sigma$ denotes the usual Haar measure on $M_3^s(\mathbb{R})$ and Z is the normalization constant. The physical parameters B, G are called the *bulk modulus* and the *shear modulus* respectively. This expression must, however, be renormalized at lower temperature, when approaching the glass transition. This renormalization, proposed by Eshelby [8], is the analog for continuum mechanics to the Clausius–Mossotti formula for dielectric. Namely the long distance propagation of stress, due to the elasticity equation, creates a local effective stress field on the local atom which can be summarized as a modification of the values of B, G. This numerical study suggests that, in the liquid phase, the stress degrees of freedom are dominant and behave like a free gas at equilibrium, in complete analogy with the Einstein theory of phonons in crystals. In particular, the heat capacity becomes elementary to compute and behaves like a constant at high temperature: this is the origin of the *Law of Dulong–Petit*, confirmed by experiments [4] (see Figure 1 and eq. (1) in Section 1.3). The main question is to give this finding an explanation from first principle.

3.2. Anankeon

The intuitive reason why the local stress is actually random in the liquid phase can be seen by imagining what the life of the local atom should be. The atom can never find a position that minimizes all pair potential energy and this is the result of the local frustration. Even in the ideal situation, for which the local cluster is made of identical atoms localized on the center and the vertices of a perfect icosahedron, the distance between the central atom to the vertex-ones is 5% less than between two vertex-atoms, meaning that it is not possible to minimize all pair potentials simultaneously. The random thermal motion then intervenes to move atoms around in search for a better position. Hence the corresponding stress

tensor felt by each atom is constantly varying under the *stress of circumstances*, a concept that can be translated by the word $\alpha\nu\alpha\gamma\kappa\epsilon\iota\alpha$ (anagkeia) in Greek.

There is a a character of the Greek mythology that could fit with this concept, the goddess *Ananke*, whose name comes from the greek word anagkeia. Ananke was representing a power above all including the Gods of the Greek Mythology. It expresses the concepts of "force, constraint, necessity" and from there it also means "fate, destiny" to lead to the concepts of compulsion, torture (see the entry Ananke in Wikipedia). From this, the word *anankeon* can be coined to describe the dominant degrees of freedom taking place in the liquid phase [1]. What are these anankeons? It is the purpose of what follows to make this concept more precise. In the physics literature focussing on this problem, these degrees of freedom are also called *topological*, meaning that the local atomic arrangement, also called local clusters, or local patches in tiling theory, is modified. Another way to put it is to describe the association of atoms in terms of *bonds* in a way similar to chemical bonds. The movement of atoms corresponds to a random change in the bonding between atoms.

4. Topology and combinatoric of the space of Delone sets

It will be argued that a Voronoi construction [44] (also called Dirichlet tesselation) on the set of configurations is the right mathematical tool liable to describe anankeons. In what follows, the *Delone set Hypothesis* is assumed to hold and all atomic configurations will be chosen in $\mathrm{Del}_{r,R}(\mathbb{R}^d)$ which will be denote by Del. For simplicity the specification of various atomic species and of their relative concentration will be omitted here. Then Del^0 will denote the set of such configurations *having one atom at the origin*. This section is a summary of a work in preparation [38]. In particular all Theorems and other claims will be proved there. The author suspects that several of these results are already known but did not find yet the corresponding references.

4.1. Delaunay triangulation

The Voronoi tiling point of view, was already used by Bernal [3]. Given a discrete Delone set $\mathcal{L} \subset \mathbb{R}^d$, define the *Voronoi cell* V_x of a point $x \in \mathcal{L}$ (called atoms here) as the set of points in the space closer to x than to any other atom in $\mathcal{L}$. Since $\mathcal{L}$ is Delone, say with parameters (r, R), it follows that V_x is an *open convex polytope* containing the open ball $B(x; r)$ and contained in the ball $B(x; R)$. The *Voronoi tile* at x is just the closure $T_x = \overline{V_x}$. Two distinct Voronoi cells have no point in common, while two distinct Voronoi tiles may intersect only on a common face. Two distinct atoms $x, y \in \mathcal{L}$ will be called *nearest neighbors* whenever their Voronoi tiles are touching along a common face of codimension one (namely of dimension $d-1$). The set $\mathcal{E}$ of pairs $e = (x, y)$ of nearest neighboring atoms, called *edge* or *bonds*, defines with $\mathcal{L}$ a graph $\mathcal{G} = (\mathcal{L}, \mathcal{E})$ which will be called the *Delone graph*. It is known in the literature as the *Delaunay Triangulation* [40]. It plays the role of a dual lattice. This graph is simple (between any two vertices there

is at most one edge), connected (any two vertices can be linked by a finite path, namely a sequence of edges touching by a common vertex). Moreover the length of any edge lies between $2r$ and $2R$. Hence the graph distance $d_{\mathcal{G}}(x, y)$ between two atoms, defined as the length of the shortest path linking x to y, is equivalent to the Euclidean distance.

4.2. Voronoi points, Delaunay triangulation and the empty sphere property

In this work the name *Voronoi point* will be given to a vertex of one of the Voronoi tiles, namely a face of dimension zero (see Figure 3). An atomic neighbor (or a-neighbor) of a Voronoi point is an atom with Voronoi tile containing this point. For example, in Figure 3, the Voronoi point indicated on the left has 4-a-neighbours, while a small deformation of the atomic positions, as seen on the right shows that this Voronoi point has split and each of the new Voronoi points has exactly 3-a-neighbors (generic situation). It is important to remark that atomic sites have a material existence in a solid, while Voronoi points are just a convenient mathematical concept that does not correspond to any material point in the solid.
The following result is classic and goes back to [40]

Theorem 2. *Let* $\mathcal{L} \in \mathrm{Del}^0$.

(i) *Any Voronoi points of* $\mathcal{L}$ *has at least* $d+1$ *a-neighbors. It belongs to the interior of the convex hull of its a-neighbors. Generically, in* Del^0*, all Voronoi points have exactly* $d+1$ *neighbors.*

(ii) **Empty sphere property:** *The set of a-neighbors of a Voronoi point* y *of* $\mathcal{L}$ *is contained in a sphere centered at* y *the interior of which contains no other atoms. Conversely, any sphere in* $\mathbb{R}^d$ *defined by at least* $d+1$ *atoms of* $\mathcal{L}$ *which does not contained atoms in its interior admits a Voronoi point at its center.*

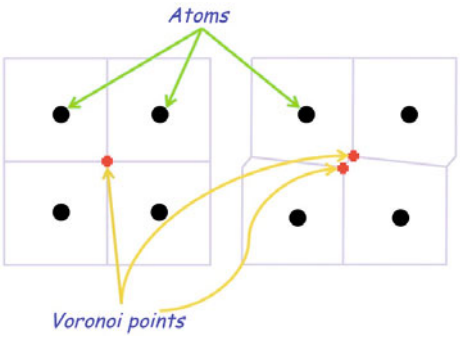

FIGURE 3

A Voronoi point will be called *simple* whenever it has exactly $d+1$ a-neighbors. Generically, all Voronoi points of $\mathcal{L}$ are simple. Hence they are at the

center of an empty sphere and the corresponding atoms give elementary d-simplex (triangles in 2D and tetrahedra in 3D) making up a triangulation of the space, called the *Delaunay triangulation.*

4.3. Graph balls and local topology

A *graph ball* is a subgraph $\mathcal{G}(x,n) = (\mathcal{V}(x,n), \mathcal{E}(x,n))$, where $\mathcal{V}(x,n)$ is the set of all atoms at a graph distance less than or equal to $n \in \mathbb{N}$ from x, and $\mathcal{E}(x,n)$ is the set of all edges $(y,z) \in \mathcal{E}$ such that $y, z \in \mathcal{V}(x,n)$. Then x is the center and n the radius of this graph ball. The *support* of a graph ball is the union of Voronoi tiles centered at its vertices.

Two graphs $\mathcal{G}, \mathcal{G}'$ are called *isomorphic* if there are bijective maps $\phi^v : \mathcal{V} \to \mathcal{V}'$ and $\phi^e : \mathcal{E} \to \mathcal{E}'$ compatible with the edge-boundary, namely the vertices linked by the image $\phi^e(e)$ (the boundary) of an edge e is the ϕ^v-image of the vertices liked by e. Equivalently a graph isomorphism is an isometry of its vertex set endowed with the graph distance.

A *local patch* is an isomorphism class of graph balls. In particular all such balls have the same radius. In this sense a *local patch is topological*, because it is likely to be insensitive to small move of the corresponding atoms. It is essential to remark that

Proposition 3. *The number of local patches of given radius is finite namely the Delone Graphs have Finite Local Complexity.*

In other words, when the atoms are moving a bit, the corresponding Delone graph is deformed, but its isomorphism class does not change.

4.4. Alloys

If $\mathcal{L}$ represent the position of the atoms in an *alloy* with more than one species, the Voronoi construction can be done provided the Euclidean distance is replaced by a local metric taking care of the different atomic radii. Namely V_x is defined as the set of points $z \in \mathbb{R}^d$ such that $\|z-x\|/a_x < \|z-y\|/a_y$ for all $y \in \mathcal{L}$ distinct from x, where a_x denotes the atomic radius of the atom x. Then the Voronoi cells are *curved polytopes*, namely their faces are made of pieces of spheres. But all the combinatorial aspect described here will persist.

4.5. Acceptance domains

Given a local patch $\mathcal{P}$ of radius n, its acceptance domain $D(\mathcal{P})$ is the set of all Delone sets in Del^0 with graph ball of radius n, centered at the origin, given by $\mathcal{P}$. Given $n \in \mathbb{N}$, the family of all acceptance domains of local patches of radius n is a *partition* of Del^0. A local patch will be called *generic* whenever its acceptance domain is open.

Theorem 3 ([38]). *A local patch $\mathcal{P}$ of radius n is generic if and only if the support of the n-graph ball centered at the origin of any atomic configuration in its acceptance domain contains only simple Voronoi points.*

The closure of an open acceptance domain will be called a *domain of contiguity.*

Theorem 4 ([38]). *The set of domains of contiguity of patches of radius n is a finite cover of* Del^0.

In a sense the family of generic patches of given radius generates a finite tiling of Del^0, if the tiles are the domain of contiguity while their interiors are the open acceptance domains. Moreover passing from n to $n+1$ gives a *refinement* of this tiling, namely any contiguity domain of radius n is tiled by the contiguity domains of radius $n+1$ it contains. The following shows that genericity coincides with full probability with respect to any Gibbs state.

Theorem 5 ([38]). *Given $\mathbb{P}$ a locally absolutely continuous probability measure on* Del^0 *the boundary of a contiguity domain has zero probability.*

4.6. Graph of contiguity

Given $n \in \mathbb{N}$, two contiguity domains of radius n are *contiguous* if their intersection contains a subset of co-dimension one. Then the *graph of contiguity* $\mathfrak{G}_n = (\mathfrak{V}_n, \mathfrak{E}_n)$ is defined as the graph with set of vertices $\mathfrak{V}_n$ given by all domain of contiguity of radius n (or equivalently by the set of generic patches) and edge set $\mathfrak{E}_n$ given by contiguous pair of domains of contiguity in $\mathfrak{V}_n$. As it turns out

Theorem 6 ([38]). *$\mathfrak{G}_n$ is a finite, simple and connected graph. The number of its vertices increases exponential fast with n.*

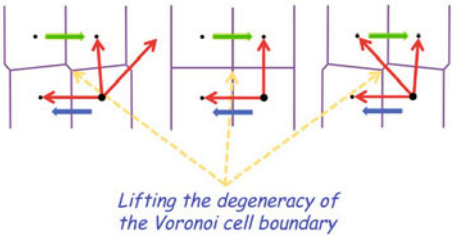

FIGURE 4. Pachner move in dimension 2: a $2 \leftrightarrow 2$ Voronoi point collision.

4.7. Contiguity and Pachner moves

The previous description makes sense from a mathematical standpoint. But it looks extremely complicate at first sight, since the space Del^0 is infinite dimensional: indeed, the restriction of Del^0 to any bounded domain U can be seen as the disjoint union of a family of open sets in $\mathbb{R}^{nd}$, where n denotes the number of atoms inside U. It is actually simple to describe a transition from a domain of contiguity to a contiguous one. For indeed, such a move requires only one local change in the Delone set. To express this fact more precisely let $n \in \mathbb{N}$ be chosen as the radius of patches. Then a *generic atomic configuration belonging to the intersection $\mathcal{T}_0 \cap \mathcal{T}_1$ of two contiguous domains of contiguity is one with only one non simple Voronoi*

point within the ball of radius n centered at the origin. In addition, this Voronoi points admits exactly $d+2$ *a-neighbors.* In particular:

(i) in dimension 2, this Voronoi point admits 4 atomic neighbors (see Figure 4),
(ii) in dimension 3 it has 5 atomic neighbors (see Figure 5).

FIGURE 5. Pachner move in dimension 3: a $2 \leftrightarrow 3$ Voronoi point collision.

As soon as the configuration moves away from the $\mathcal{T}_0 \cap \mathcal{T}_1$, this nongeneric Voronoi point splits into a family of generic ones and this splitting is known in triangulation theory used in Geometry as a *Pachner move* [43] or *bistellar flips*. More precisely (see Figure 4 and Figure 5),

(a) in dimension 2, two generic Voronoi points collide at the transition to give two other ones;
(b) in dimension 3, two generic Voronoi points collide to give three other ones or vice-versa;
(c) in higher dimension l Voronoi points collide at the transition to give $m = d+2-l$ other ones for $2 \leq l \leq d$.

5. Applications to material science

The previous mathematical program can be applied in practical circumstances. Pachner moves, seen as collisions of Voronoi points leads to several practical consequences. In a sense the anankeon can be seen as the degree of freedom associated with the motion of Delaunay simplex,(triangle in dimension 2, tetrahedron in dimension 3) formed by the $d+1$ atoms around the Voronoi point. Hence an anankeon can be labeled by this Voronoi point.

5.1. Pachner moves in materials

The previous results should apply in real material, apart from taking energy considerations into account. A Pachner move corresponds to a mechanical motion of the atomic configuration passing through a *saddle point* of the potential energy. In particular if the atomic configuration corresponds to a point of the energy surface close enough to the saddle point, it looks almost like the perfect nongeneric Voronoi point. In practice it means that when several Voronoi points are very

close to each other, then the configuration is physically equivalent to a nongeneric unique Voronoi point. In the scheme described in Section 4.5, it means that the boundary of a domain of continuity is actually thicken by the energetic considerations. For this reason the number of atoms involved in a Pachner move is likely to be a random variable. This is what has been observed recently in the Egami group [49] (see Figure 6).

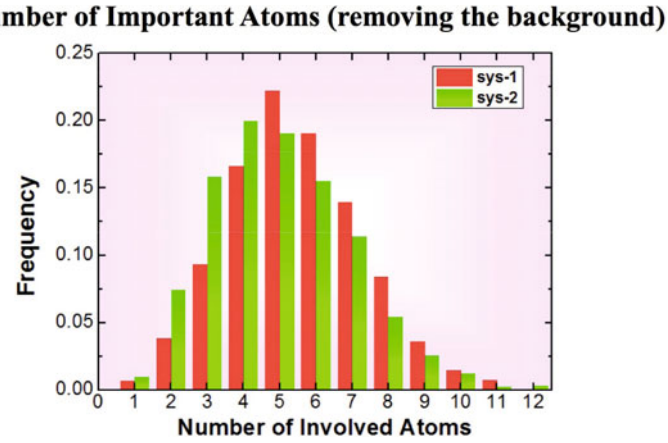

FIGURE 6. Distribution of the number of atoms involved in a Pachner move [49].

It is remarkable that this distribution is Poissonian with average 5 as predicted by the Pachner theory in dimension 3.

5.2. Pachner moves and STZ

The *shear transformation zone* theory (STZ) proposed by Langer *et al.* [9, 10, 15, 16, 17], consists simply in remarking that indeed such Pachner moves occur in glasses when submitted to a strong shear. In the glass phase, however, the atoms are blocked in their position, so that the dominant degree of freedom is provided by the phonons only. However, if a shear is applied, first the solid reacts elastically. But beyond some critical value of the shear, the deformation becomes irreversible. At the atomic scale, this can be interpreted as one Pachner move somewhere and the numerical simulations show that such a move does not come alone, it is followed by a cascade of moves in the vicinity. In other words a small region about the germ behaves like a liquid, and this is exactly what an STZ is. The spatial distribution of such STZ is random, presumably Poissonian with a density depending upon the applied shear. This is exactly the hypothesis made in the STZ theory. Each Pachner moved is *polarized*, as can be figured out from Figures 4 and 5. In 2D a Pachner moves corresponds the a switch of a diagonal bond from left to right or vice-versa. In 3D it corresponds to a collision of 2 Voronoi points giving 3 or vice-versa. In both cases, there is a change between two types of configurations,

each of which being called a *polarization*. Then the local parameter is provided by the numbers $N_\pm$ of each of these varieties of groups Voronoi points in the STZ. This number is a random variable the time evolution of which is assumed to be Markovian, namely it follows a Master Equation. As it turns out, this effect can be added to the equation of elasticity, including the temperature effects. Eliminating the parameters $N_\pm$ from them leads to a non linear partial differential equation which describe accurately the plasticity region in the strain-stress relation for a metallic glass [10]. Using this equation it becomes possible to simulate numerically in a realistic way a fracture of the material [20].

5.3. Atomic scale Markov process

In order to describe the anankeon dynamics, the previous considerations leads to represent it by a Markov process on the graph of contiguity $\mathfrak{G}_n$. The radius n is simply a measure of the volume of the region under consideration. A Markov process represents the dynamics of a random walker moving on the set of vertices of the graph through the edges (standard references are [47, 48] for discrete time and [45, 46] for continuous time). Since $\mathfrak{G}_n$ is a finite graph, a Markov process can be equivalently described by its generator $\mathfrak{L}_n$, namely a linear operator acting of the linear space $\mathcal{H}_n$ of complex-valued functions defined on $\mathfrak{V}_n$, as follows

$$\mathfrak{L}_n(f)(\mathcal{P}) = \frac{1}{T_\mathcal{P}} \sum_{\mathcal{Q}\sim\mathcal{P}} \mathbb{P}_{\mathcal{Q}\leftarrow\mathcal{P}} \left(f(\mathcal{P}) - f(\mathcal{Q})\right) .$$

in such an expression, $\mathcal{P}, \mathcal{Q}$ label patches of radius n which are contiguous, and $\mathbb{P}_{\mathcal{Q}\leftarrow\mathcal{P}}$ represent the probability to jump from $\mathcal{P}$ to $\mathcal{Q}$. Moreover $T_\mathcal{P}$ represents the *waiting time* of the random walker at the patch $\mathcal{P}$. Since this Markov process is supposed to define the thermal equilibrium motion, the jump probability is given by a Gibbs factor of the from

$$\mathbb{P}_{\mathcal{Q}\leftarrow\mathcal{P}} = \frac{1}{\mathcal{Z}} e^{-\beta(F(\mathcal{Q})-F(\mathcal{P}))} ,$$

where $F(\mathcal{P})$ represents the *free energy* of the configurations represented by $\mathcal{P}$, while $\mathcal{Z}$ is a normalization constant ensuring that the sum over $\mathcal{Q}$ of the *r.h.s.* equals 1. On the other hand, the waiting time is provided by an Arrhenius law and depends only of the potential barrier $W(\mathcal{P}\to\mathcal{Q})$ between the initial configuration and the saddle points during Pachner moves

$$T(\mathcal{P}) = e^{\beta W(\mathcal{P}\to\mathcal{Q})} .$$

Numerical simulations based on molecular dynamics permit to estimate precisely the previous parameters [41, 42] and, in principle, to test the validity of such an approach. However, for the previous description to be realistic, the atomic stress degrees of freedom must be introduced in the Markov modeling.

5.4. Viscosity and glass transition

Once the Markov dynamics is defined and tested, it becomes possible to compute the viscosity. It is defined as the linear response to an applied shear in terms of

the time derivative of the strain. Thanks to the Green–Kubo formula, it is also equal to the stress-stress correlation function. As the temperature decreases, the thermal activation moving the atoms around becomes less effective. In particular given some shear, it takes more Pachner moves to achieve the same amount of strain variation. Hence the typical time scale involved in the viscosity (called the Maxwell time) becomes much larger than the time scale involved in the definition of the Markov process. The glass transition is achieved when the Maxwell time becomes so large as to make it impossible to observe. Using this principle it becomes possible to give a formula for the temperature of the glass transition [39] fitting with experimental datas.

6. Geometry on the space of Delone configurations

This Section is conjectural and proposes a program liable to define on the compact space Del^0 a *geometry*, namely a metric, liable to describe, in analogy with General Relativity, the dynamics of a condensed material in terms of a geodesic flow with respect to this metric structure.

6.1. Infinite volume limit

In Section 5.3 a Markov process is proposed as a modeling of the time evolution of a liquid or glassy material. This Markov process, though, is defined on graph balls of a fixed finite radius $n \in \mathbb{N}$. The question remains whether the sequence of such process converges to some limit in the limit where the radius n tends to infinity. This is an entirely open problem so far. However, there are indications that it could be possible to prove the convergence. One strategy could be to use the martingale convergence Theorem [57, 58]. It is a technique used elsewhere to prove the existence of a diffusion process on fractals like the Sierpinski carpet [50]. If so, then the generator $\mathfrak{L}$ of this limiting Markov process will acts on the entire Del^0 space as a Dirichlet form [51, 52, 59]. Another strategy could be to use convergence Theorem for Markov semigroups on C^*-algebras [53].

6.2. Dirichlet forms and gradient operators

It is important though that the Markov process has the so-called *Feller property*, namely that the image of a continuous function on Del^0 by the process stays continuous. If so, this Markov process should be represented as a Markov semigroup on the C^*-algebra $\mathcal{A} = \mathcal{C}(\mathrm{Del}^0)$, namely the space of continuous functions on Del^0. More is needed though: namely *the domain of* $\mathfrak{L}$ *should contain a dense* $*$*-subalgebra* of $\mathcal{A}$. In a such a case, an extension of the Lindblad Theorem leads to the following: there are

(i) a C^*-right $\mathcal{A}$-module $\mathcal{E}$
(ii) a $*$-representation $\pi : \mathcal{A} \to \mathrm{End}(\mathcal{E})$
(iii) a linear map $\partial : \mathcal{C}(\mathrm{Del}^0) \to \mathcal{E}$ such that, for every $f, g \in \mathcal{A}$

$$\mathfrak{L}(f^* g) - f^* \mathfrak{L}(g) - \mathfrak{L}(f^*) g = \langle \partial f | \partial g \rangle_{\mathcal{E}}$$

and that ∂ becomes a derivation, namely

$$\partial(fg) = \partial(f)g + \pi(f)\partial(g)\,.$$

In other words, ∂ is a candidate for a generalized version of the *gradient* operator. Since the inner product $\langle\,\cdot\,|\,\cdot\,\rangle_{\mathcal{E}}$ takes on values in $\mathcal{A} = \mathcal{C}(\mathrm{Del}^0)$, it follows that

$$\|f\|_{Lip} \stackrel{\mathrm{def}}{=} \sup_{\nu\in\mathrm{Del}^0} |\langle\partial f|\partial f\rangle_{\mathcal{E}}(\nu)|^{1/2}\,,$$

plays the role of a a Lipshitz norm. Hence, by the Kantorovich duality [61, 62] theory, it is likely to define a *metric* on Del^0 through the formula

$$d_{\mathfrak{L}}(\nu,\nu') = \sup\{|f(\nu) - f(\nu')|\,;\,\|f\|_{Lip} \le 1\}\,.$$

Problem 3. Prove that all conditions are fulfilled to make sure that the metric $d_{\mathfrak{L}}$ is well defined and generates the weak*-topology on Del^0. □

Such a problem has already been investigated in full generality by Rieffel in the context of *Quantum Metric Spaces* [65, 66, 67, 68, 63], providing a criterion to answer positively to this question.

6.3. The curvature problem

The previous construction, if it works, leads to a structure of compact metric space on Del^0 generated by the dynamics describing the time evolution for glasses and liquids. One important question is to check whether such a structure share with Riemannian manifolds a notion of *curvature*. In Noncommutative Geometry, thanks to the work of I. Palmer [64], a Riemannian structure can be defined on any compact metric space using the concept of *spectral triple* [55]. However since the Hausdorff dimension of Del^0 is infinite, the construction cannot be used very far. In particular it is likely that a *Connes state*, namely a Dixmier trace, does not exist. Hence one can only expect a KMS state with respect to the Dirac operator associated with the spectral triple.

Problem 4. Elucidate the meaning of the KMS condition produced in this way. Can one see the Connes state as the equilibrium invariant state associated with the Markov dynamics generated by $\mathfrak{L}$? □

In addition, the definition of the *curvature* has not yet reached a consensus in Noncommutative Geometry [56]. So, it might be necessary to come back to the usual approach in *Metric Geometry* [60, 54]. The celebrated Gromov Theorem asserts that the set of Riemannian manifolds with uniformly bounded volume, curvature and injectivity radius, is precompact in the Hausdorff–Gromov metric topology. It raises the question of what is a limit point in this topology. Examples are given by *Alexandrov spaces*, namely path-metric spaces in which geodesic exist between any pair of points, the concept of angle between two geodesic starting at the same point is well defined, and in which a concept of scalar curvature exists [54]. The curvature is obtained from comparing geodesic triangles in this space with isometric ones in model Riemannian manifolds with constant curvature [54].

Problem 5. If the construction of Section 6.2 can be made, is it possible to see Del^0 as an Alexandrov space? If yes, is there an interpretation of the curvature in terms of the Continuum Mechanics of the condensed material under scrutiny? □

Such a relation is exactly given by the Einstein equation in General relativity, linking the gravitational force to the Geometry of the Universe. This series of question suggests to look at whether such an interpretation is likely to hold in the description of a condensed material.

Acknowledgment

The author is especially indebted with T. Egami for many exchanges of idea since February 2012. He is also indebted with J. Langer who accepted to have a long discussion in May 2013 about his ideas on the STZ theory.

References

Introduction

[1] The name *anankeon* used in these notes is the result of an exchange of emails between Takeshi Egami and the author in Fall 2012. It has been used for the first time in a publication in [7].

[2] J. Bellissard, D. Hermmann, M. Zarrouati, "Hull of Aperiodic Solids and Gap labeling Theorems", in *Directions in Mathematical Quasicrystals*, CRM Monograph Series, **13**, 207–259, M.B. Baake & R.V. Moody, eds., AMS Providence, (2000).

[3] J.D. Bernal, "The Bakerian Lecture, 1962. The Structure of Liquids", *Proc. R. Soc. A*, **280**, (1984), 299–322, and references therein.

[4] H.S. Chen, D. Turnbull, "Evidence of a Glass-Liquid Transition in a Gold-Germanium-Silicon Alloy", *J. Chem. Phys.*, **48**, 2560–2571, (1968).

[5] S.-P. Chen, T. Egami, V. Vitek, *Phys. Rev. B*, **37**, (1988), 2440–2449.

[6] T. Egami, Y. Waseda, "Atomic size effect on the formability of metallic glasses", *J Non-Cryst Solids*, **64**, (1984), 113–134.

[7] T. Egami, "Elementary Excitation and Energy Landscape in Simple Liquids", *Mod. Phys. Lett. B*, **28**, (2014), 1430006 1–19.

[8] J.D. Eshelby, "The Determination of the Elastic Field of an Ellipsoidal Inclusion, and Related Problems", *Proc. R. Soc. London*, **241**, (1957), 376–396.

[9] M.L. Falk, J.S. Langer, "Dynamics of viscoplastic deformation in amorphous solids", *Phys. Rev. E*, **57**, (1998), 7192–7205.

[10] M.L. Falk, J.S. Langer, L. Pechenik, "Thermal effects in the shear-transformation-zone theory of amorphous plasticity: Comparisons to metallic glass data", *Phys. Rev. E*, **70**, (2004), 011507.

[11] J.P. Hansen, I.R. McDonald, *Theory of Simple Liquids*, Second Edition, Academic Press, London (1990).

[12] A. Inoue, "Stabilization of metallic supercooled liquid and bulk amorphous alloys", *Acta Mater.*, **48**, (2000), 279–306.

[13] W.L. Johnson, "Bulk Glass-Forming Metallic Alloys: Science and Technology", (1998 Material Research Society Medal Award Lecture, Presented at Symposium), in *MRS*

Proceedings, Volume **554**, (1998); "Bulk glass-forming metallic alloys: science and technology", *Mater. Res. Soc. Bull.*, **24**, (1999), 42–56.

[14] L.D. Landau, E.M. Lifshitz, *Theory of Elasticity*, Volume 7 of Course on Theoretical Physics, 3rd Edition, reviewed and enlarged by E.M. Lifshitz, A.M. Kosevitch and L.P. Pitaevskii,1986, Butterworth-Heinemann Linacre House Pub., Oxford, 1999.

[15] J.S. Langer, L. Pechenik, "Dynamics of shear-transformation zones in amorphous plasticity: Energetic constraints in a minimal theory", *Phys. Rev. E*, **68**, (2003), 061507.

[16] J.S. Langer, "Dynamics of shear-transformation zones in amorphous plasticity: Formulation in terms of an effective disorder temperature", *Phys. Rev. E*, **70**, (2004), 041502.

[17] J.S. Langer, "Shear-transformation-zone theory of deformation in metallic glasses", *Scripta Materialia*, **54**, (2006), 375–379.

[18] D.B. Miracle, "A structural model for metallic glasses", *Nature Mater.*, **3**, 697–702, (2004).

[19] C.H. Rycroft, E. Bouchbinder, "Fracture Toughness of Metallic Glasses: Annealing-Induced Embrittlement", *Phys. Rev. Letters*, **109**, (2012), 194301:1–5.

[20] C.H. Rycroft, F. Gibou, "Simulations of a stretching bar using plasticity model from the shear transformation zone theory", *J. Comp. Phys.*, **231**, (2012), 2155–2179.

[21] S.A. Silling, "Reformulation of elasticity theory for discontinuities and long-range forces", *J. Mech. Phys. Solids*, **48**, (2000), 175–209.

[22] S.A. Silling, , M. Epton, O. Weckner, J. Xiu, E. Askari, "'Peridynamics states and constitutive modeling', *J. Elasticity*, **88**, (2007), 151–184.

[23] S.A. Silling, R.B. Lehoucq, "Convergence of Peridynamics to Classical Elasticity Theory", *J. Elasticity*, **93**, (2008), 13–37.

Why Delone sets?

[24] J. Bellissard, C. Radin, S. Shlosman, "The characterization of ground states", *J. Phys. A: Math. Theor.*, **43**, (2010), 305001.

[25] C.H. Bennett, P. Chaudhari, P. Moruzzi, P.J. Steinhardt, "On the stability of vacancy and vacancy clusters in amorphous solids", *Phil. Mag.* **A40**, (1979), 485–495.

[26] R.L. Dobrushin, "Gibbsian Random Fields for Lattice Systems with Pair Interactions", *Funct. Anal. Appl.*, **2**, (1968), 291–301.

[27] O.E. Lanford III, Entropy and equilibrium states in classical Statistical Mechanics", in *Statistical Mechanics and Mathematical Problems*, Lecture Notes in Physics Volume **20**, 1973, pp 1–113.

[28] O.E. Lanford III, "Time evolution of large classical systems", in *Dynamical systems, theory and application*, Lecture Notes in Physics Volume **38**, 1974, pp. 1–111.

[29] S. Majumdar, "Persistence in Nonequilibrium Systems", *Curr. Sci.*, **77**, (1999), 370–375, accessible on `arXiv cond-mat/9907407`.

[30] D. Ruelle, *David Statistical mechanics: Rigorous results*, W. A. Benjamin, Inc., New York-Amsterdam (1969).

[31] D. Ruelle, "Superstable interactions in classical statistical mechanics", *Comm. Math. Phys.*, **18**, (1970), 127–159.

The anankeon theory

[32] R. Benguria, E.H. Lieb, "The positivity of the pressure in Thomas–Fermi theory", *Comm. Math. Phys.*, **63**, 193–218.

[33] T. Egami, D. Srolovitz, "Local structural fluctuations in amorphous and liquid metals: a simple theory of the glass transition", *J. Phys. F: Met. Phys.*, **12**, (1982), 2141–2163.

[34] E.H. Lieb, "Thomas–Fermi and related theories of atoms and molecules", *Rev. Modern Phys.*, **53**, (1981), 603–641; "Erratum", *Rev. Modern Phys.*, **54**, (1982), 311.

[35] M. Mézard, G. Parisi, M.A. Virasoro, *Spin glass theory and beyond*, World Scientific, Singapore, (1987).

[36] K.R. Parthasarathy, *Probability measures on metric spaces*, American Mathematical Soc., 1967.

[37] W.H. Wang, C. Dong, C.H. Shek, "Bulk metallic glasses", *Mat. Sc. Engin. R*, **44**, (2004), 45–89.

Topology and combinatoric of the space of Delone sets

[38] J. Bellissard, paper in preparation.

[39] T. Egami, S.J. Poon, Z. Zhang, V. Keppens, "Glass transition in metallic glasses: A microscopic model of topological fluctuations in the bonding network", *Phys. Rev. B*, **76**, (2007), 024203, 1–6.

[40] B. Delaunay, "Sur la sphère vide. A la mémoire de Georges Voronoï", *Bull Acad. Sci. URSS*, **6**, (1934), 793–800.

[41] T. Iwashita, T. Egami, "Atomic mechanism of flow in simple liquids under shear", *Phys. Rev. Lett.*, **108**, (2012), 196001 1–4.

[42] T. Iwashita, D.M. Nicholson, T. Egami, "Elementary Excitations and Crossover Phenomenon in Liquids", *Phys. Rev. Lett.*, **110**, (2013), 205504, 1–5.

[43] U. Pachner, "P. L. homeomorphic manifolds are equivalent by elementary shellings", *European J. Combin.*, **12**, (1991), 129–145.

[44] G. Voronoi, "Nouvelles applications des paramètres continus à la théorie des formes quadratiques", *J. Reine Angew. Math.*, **133**, (1908), 97–178.

Applications to material science

[45] R.M. Blumenthal, R.K. Getoor, *Markov processes and potential theory*, Monographs in Pure and Applied Mathematics, Academic Press, New-York, (1968), reprint by Elsevier, (2011).

[46] E. Çinlar, *Probability and Stochastics*, see Chapt. IX, Springer, (2011).

[47] J.R. Norris, *Markov Chains*, Cambridge University Press, (1998).

[48] D.W. Stroock, *An Introduction to Markov Processes*, Springer, (2005).

[49] Y. Fan, T. Iwashita, T. Egami, "How Thermally Induced Deformation Starts in Metallic Glass", *Nature Communications*, in press (2014).

Geometry on the space of Delone configurations

[50] M.T. Barlow, R.F. Bass, "The construction of Brownian motion on the Sierpinski carpet", *Ann. I.H.P.*, **25**, (1989), 225–257.

[51] A. Beurling, J. Deny, "Espaces de Dirichlet. I. Le cas élémentaire", (French) *Acta Math.*, **99**, (1958), 203–224.

[52] A. Beurling, J. Deny, "Dirichlet spaces", *Proc. Nat. Acad. Sci. U.S.A.*, **45**, (1959), 208–215.

[53] O. Bratteli, D.W. Robinson, *Operator Algebras and Quantum Statistical Mechanics*, Vol. 2, Springer-Verlag, 2nd ed. (1997).

[54] D. Burago, Y. Burago, S. Ivanov, *A Course in Metric Geometry*, Graduate Studies in Mathematics, 33. American Mathematical Society, Providence, RI, 2001.

[55] A. Connes, *Noncommutative Geometry*, Academic Press, (1995).

[56] A. Connes, H. Moscovici, "Modular Curvature for Noncommutative Two-Tori", `arXiv:1110.3500`, (2011).

[57] J.L. Doob, "Regularity properties of certain families of chance variables", *Trans. Amer. Math. Soc.*, **47**, (1940), 455–486.

[58] J.L. Doob, *Stochastic processes*, John Wiley & Sons, Inc., New York; Chapman & Hall, Limited, London, (1953).

[59] M. Fukushima, *Dirichlet Forms and Markov Processes*, North-Holland, (1980).

[60] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Birkhäuser, Boston (1999).

[61] L.V. Kantorovich, "On one effective method of solving certain classes of extremal problems", *Dokl. Akad. Nauk USSR* **28** (1940), 212–215.

[62] L.V. Kantorovič, G. Š. Rubinšteĭn, "On a functional space and certain extremum problems", (Russian) *Dokl. Akad. Nauk SSSR* (*N.S.*) **115** (1957), 1058–1061; "On a space of completely additive functions", (Russian) *Vestnik Leningrad. Univ.* **13** (1958), 52–59.

[63] N. Ozawa, M. Rieffel, "Hyperbolic group C^*-algebras and free-product C^*-algebras as compact quantum metric spaces", *Canad. J. Math.* **57** (2005), no. 5, 1056–1079.

[64] I.C. Palmer, *Noncommutative Geometry and Compact Metric Spaces*, PhD Thesis, Georgia Institute of Technology, (May 2010).

[65] B. Pavlović, "Defining metric spaces via operators from unital C^*-algebras", *Pacific J. Math.* **186** (1998), 285–313.

[66] M. Rieffel, "Metrics on states from actions of compact groups", *Doc. Math.* **3** (1998), 215–229.

[67] M. Rieffel, "Metrics on states spaces", *Doc. Math.* **4** (1999), 559–600.

[68] M. Rieffel, "Compact quantum metrics spaces", in *Operator Algebras, Quantization, and Noncommutatve Geometry: a Centennial Celebration Honoring John von Neumann and Marshall H. Stone*, (Doran, R.S. and Kadison, R.V., eds.), vol. 365 of *Contemporary Mathematics*, AMS, 2004, pp. 315–330.

Jean V. Bellissard
Georgia Institute of Technology
School of Mathematics
Atlanta, GA 30332-0160, USA
e-mail: jeanbel@math.gatech.edu